



Wiley Trading

TRADING REGIME ANALYSIS

The
Probability
of Volatility

MURRAY GUNN

TRADING REGIME ANALYSIS

The Probability of Volatility

Murray Gunn



A John Wiley and Sons, Ltd, Publication

Trading Regime Analysis

For other titles in the Wiley Trading Series
please see www.wiley.com/finance

Copyright © 2009

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester,
West Sussex PO19 8SQ, England

Telephone (+44) 1243 779777

Email (for orders and customer service enquiries): cs-books@wiley.co.uk

Visit our Home Page on www.wiley.com

All Rights Reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, Saffron House, 6-10 Kirby Street, London, EC1N 8TS, UK, without the permission in writing of the Publisher. Requests to the Publisher should be addressed to the Permissions Department, John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England, or emailed to permreq@wiley.co.uk, or faxed to (+44) 1243 770620.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The Publisher is not associated with any product or vendor mentioned in this book.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the Publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Other Wiley Editorial Offices

John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA

Jossey-Bass, 989 Market Street, San Francisco, CA 94103-1741, USA

Wiley-VCH Verlag GmbH, Boschstr. 12, D-69469 Weinheim, Germany

John Wiley & Sons Australia Ltd, 42 McDougall Street, Milton, Queensland 4064, Australia

John Wiley & Sons (Asia) Pte Ltd, 2 Clementi Loop #02-01, Jin Xing Distripark, Singapore 129809

John Wiley & Sons Canada Ltd, 6045 Freemont Blvd, Mississauga, ONT, L5R 4J3, Canada

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Library of Congress Cataloging in Publication Data

Gunn, Murray.

Trading regime analysis : the probability of volatility / Murray Gunn.

p. cm.

Includes bibliographical references and index.

ISBN 978-0-470-98785-8 (cloth)

1. Stocks. 2. Capital market. 3. Investments. I. Title.

HG4661.G86 2009

332.63'2042—dc22

2008052049

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN 978-0-470-98785-8 (HB)

Typeset in 10/12pt Times by Integra Software Services Pvt. Ltd, Pondicherry, India

Printed and bound in Great Britain by TJ International Ltd, Padstow, Cornwall, UK

For Flora

Contents

Foreword	ix
Acknowledgements	xiii
PART I: SUPPLY AND DEMAND	1
1 There is NO Holy Grail	3
2 The “Nature” of Markets	19
3 Volatility Defined	43
PART II: EXISTING TRADING REGIME ANALYSIS	77
4 Orthodox Pattern Recognition	79
5 Japanese Candlesticks	129
6 Volume Considerations	143
7 Previous Highs and Lows	151
8 Elliott Wave Principle	161
9 Moving Average Envelopes	179
10 Bollinger Band Width	185
11 The ADX	193
12 Point and Figure Charts	199

13	Rate of Change and Divergence	207
14	Williams %R	219
15	Donchian Channels	225
16	A Nod to the Quants	239
PART III: FURTHER IDEAS FOR TRADING REGIME ANALYSIS		251
17	Implied Volatility Curves	253
18	The Volatility Smile	261
19	My MATE	275
20	Trend-Following Performance Indicator	281
21	Trading Regime Indicator	291
PART IV: COMBINING AND USING TRADING REGIME ANALYSIS		297
22	An Eclectic Approach	299
23	Applications for Traders and Investors	309
24	Trading Regime Analysis for Economists and Fundamentalists	359
25	Case Studies	365
26	There is Still No Holy Grail	389
Appendix 1: Time Fractals and the Supply/Demand Index		395
Appendix 2: Why Do Trend Lines Work?		407
Appendix 3: Examples of Trend Lines		411
References		417
Index		419

is passionate about market analysis, writing a book about it is a labour of love rather than of profit.

Secondly, this book is not intended to be a “how to trade the markets” type of book. If I had all the answers to the questions that the markets pose, or if I had a method or system that added value in every single time period, I would certainly not be writing a book about it! I would be sitting in the sunshine by the pool, at that villa by the golf course where I would play every day without having to worry about where my daughter’s education fund was going to come from! Ah yes, we can but dream. If someone has a method or arbitrage system that is so good, why would he want to share it with anyone else in the first place since doing so may well dilute its potency? No, this book is not a “how to” book in the genre of making you a guaranteed millionaire, so if that is what the reader is expecting then I suggest he puts it down now. This book is the amalgamation of my thoughts and experiences gained over the last 20 years in market trading and analysis. It is intended to be a contribution to the body of knowledge in market analysis, in particular technical market analysis that the International Federation of Technical Analysts (IFTA) is doing so much to categorise and nurture. Trading ideas and methods *are* presented in the book but they come with a very big health warning because I know from my own experience of investing in and trading the markets that no method is infallible. What we deal with in market analysis is probabilities rather than certainties. There are no certainties in the markets. Ever. This is a realisation that can really only come through market experience and having been knocked down countless times just when you think you have got the markets figured out. Probability can be measured in many different ways and it can mean different things to different people but, essentially, it is defined as the chance of something happening. In the markets that chance is generally 50% but there are times when the odds improve somewhat to skew it above the classic chance of heads or tails in a long game of coin toss. This calculation of chance has to be based on an analysis of the past history because there is, in the end, nothing else on which to base it. Whether you decide to use economic data, corporate statistics, market price action or anything else, the one common thread in market analysis is that we are all trying to find things that had a relationship with how the markets subsequently behaved in the past and, therefore, in all probability, will have that same relationship in the future.

The other important point to note is that this is not a book about quantitative analysis, and if you are a “quant” who is expecting this book to be filled with equations and abstract terms please do yourself a favour and put the book down. The title of this book refers to trading regime analysis, but I must point out from the outset that *my* concept of trading regime is, if not too different from, certainly much simpler than, the concept of regime that our quantitative and econometric friends think of, which include references to Markowitz mean variance, Markov chains, vectors, modes, state space or any other such terms that anyone without a PhD in physics would struggle to understand! I have in fact included a passing reference to what I would describe as the more quantitative analysis of trading regime, but I

Foreword

I have always been very suspicious of people who write books on financial market trading or investing. The plethora of books, such as “how to trade this way” or “how to make a million dollars from the market”, has always had me asking the same question. If the analysis, system or method that these authors present to the world are so good and consistent in extracting value (i.e. money) from the financial markets, then why oh why are they not sitting on a beach somewhere watching the profits roll in rather than writing books about it. It’s a well-known fact that a lot of authors on investment or trading in the financial and commodity markets make their living from writing about how to do it, rather than doing it themselves. Many of these authors in fact have never traded the markets but instead do a great job in marketing themselves as gurus and even brands that people who do actually trade the markets follow, sometimes on what seems like a religious basis. The old phrase is that “those that *can*, **do**, and those that *cannot*, **teach**” and the cynics among us would point out that this is true in just about any field or profession. The problem with books being written about trading in the financial markets is that the people who would be interested in them tend to analyse things in a black and white fashion given the nature of their endeavour. People who are interested in trading the markets for profit know through experience that it is one of the most difficult endeavours that human beings can subject themselves to and so once they have been doing it for a while they tend to view books on “how to trade the markets” with a very healthy dose of salt. Experienced practitioners know that there is no one method or system that is the answer to the markets puzzle and that only diligent hard work and analysis is what will pay off over the long term.

So what is different about this book? Well, the first, and most important, thing to point out to the reader is that the royalties from the sale of this book are being paid directly to a charity, so it should be crystal clear that this book has not been written for any personal monetary gain. This may seem strange to some but, as someone who

must emphasise that the regime or modes that Hamilton and others refer to in the quantitative or econometric field are different from my definition of trading regime. Most relevantly, I limit my definition of trading regime to essentially two so-called states whereas a lot of the traditional quantitative references to regime switching take multiple (macro-economic) regimes into account.

In many respects this book is intended to be an antidote to the complication and extreme mathematical models that purvey the field of investment management these days. This is not to say that there is not a place for such modelling but, having read the social science of economics by training and being a market price analyst in practice, I look at the markets more from the perspective of an artist rather than an engineer. To me, the financial and commodity markets are a never-ending experiment in *social psychology* rather than an experiment in a physics laboratory. This book, therefore, is much more market philosophy than market science.

Acknowledgements

I could not have written this book if it were not for the help and support of many people, some of whom are probably not aware that they have helped me to compile my thoughts in such a way.

I would like to thank my colleagues and professional contacts with whom I have engaged in stimulating discussions about market theory and analysis throughout the years I have been working in the financial markets. Many of these people, although holding very different viewpoints from my own, hold the same passion about the markets that I have, and so debating over market philosophy has been (and still is) a great joy. In particular, my thanks go to Kurt Magnus (the Steve Irwin of the FX market), David O'Loan (for putting the idea in my head), Tony Plummer, Ian Williams, Mark Johns, Jeremy Fand, Gerry Celaya, Tim McCullough and Ryan Shea. There are countless others whose opinions about market philosophy I value and they know who they are.

Peter Eggleston and Theodore Chen from the Royal Bank of Scotland were vital in helping me with Chapter 16 and I would like to thank them for that, as well as providing some magnificent research on the subject matter in general. Monica Greer pointed me in the right direction for further background reading on psychology which was very useful. I would like to thank Mansoor Mohi-Uddin from UBS who is, in my opinion, one of the more original and insightful analysts in the markets. The UK Society of Technical Analysts has been a constant source of inspiration for me and I would like to acknowledge the fine work of the committee under the chairmanship of Adam Sorab in recent years. The global family, which is the International Federation of Technical Analysts, is filled with exceptional people who are passionate about market analysis and I salute them all here.

I have been heavily influenced in my own market thinking by the work and writings of Robert Prechter, Ed Seykota, Adam Smith, George Soros, Ludwig von Mises, Jesse Livermore and John Henry. They are constant companions in my study.

I have been fortunate in having a close family who have encouraged me throughout the years and I would especially like to thank my mother for instilling a proactive and positive attitude in me.

Lastly, I would like to thank my wife Nicola. Were it not for her consistent encouragement and support I would not have started this project in the first place and, for that, I am deeply grateful. Writing a book is a very time-consuming endeavour and, when family time is all that really matters in this world, having a spouse who understands and encourages is most appreciated. Nicola helped directly with the compilation of the book and did everything she could to make it easy for me to spend the necessary hours. She is my soul mate, and my love for her is in a perennial bull market.

Part I

Supply and Demand

1

There is NO Holy Grail

“Everything works and nothing works.”

Richard Russell

A MARKET JOURNEY

I started my career in the financial markets the way most people do. I went to university, earned an honours degree in economics and entered my first full-time job at the age of 22 believing that the markets obeyed the laws of macro- and micro-economics I had just spent the last four years studying. How wrong I was!

I have subsequently come to realise, as I think many people do, that the financial markets do not actually behave according to the textbook laws of economics and capital market theory. In fact, my own personal belief now is that the markets can (and will!) do anything at any point of time and “prediction”, as we commonly think of it, is a totally impossible and futile exercise. However, I do not look back and think that my years studying messrs Fischer, Begg, Dornbusch, the theory of comparative advantage, the theory of purchasing power parity, the capital asset pricing model and a host of other “laws” was a complete waste of time. Not at all and I will tell you why.

There are two broad schools of thought for analysing the financial markets. One is the fundamental market analysis approach that makes use of “fundamental” data such as macro-economic statistics, corporate balance sheets and corporate profitability in order to predict the future course of prices in the markets. These statistics, backed up by perceptions and guesses about what they will be like in the future, are the so-called fundamentals of financial and commodity markets. The other approach is the technical market analysis approach that makes use of past market price data and sentiment indicators in order to anticipate the future course of prices in the markets. Proponents of the technical approach shun the fundamentals, as being irrelevant to an analysis of the actual supply and demand of the market itself that they say will be apparent in the market price alone. The technical school also differs markedly from

the fundamental school because the technical school has a much more pragmatic approach to “forecasting” future market prices than the fundamental school, which is something we will look at in detail later.

These are the two broad schools of market analysis and market participants generally fall into one group or the other for their preferred method of analysing the markets with some market participants using only one form of analysis and some using both simultaneously. However, no matter what your chosen school of analysis is, in my opinion it should be every market participant’s goal (indeed obligation) to study the methods and theories involved in each school of thought in order to gain a complete and as full an understanding of market theory as possible. I find it incredible that, even now, technical market analysis is not as widely taught in universities or finance courses as, in my opinion, it should be. The behavioural finance aspects of markets are thankfully becoming much more popular in courses though and, when you consider that behavioural finance is merely another term for technical market analysis (because technical analysis is an analysis of market psychology), then at least there is an element of balance appearing in modern finance teachings. Would we expect a lawyer specialising in criminal law only to have studied criminal law and nothing else? Would we expect a doctor specialising in cardiology only to have studied heart-related medicine? Of course not. We expect other professionals to have studied and have a broad understanding of their subject in order that we, as customers of their service, can have confidence in their knowledge and experience. Why does it seem different in the investment field? Both fundamental and technical analyses have recognised industry-accepted qualifications, but still only the minority of people in the investment industry will have badges from both schools. In this regard, therefore, my early market years submerged in the fundamental market approach have proven to be very useful indeed as I believe it has allowed me, as a keen student of the markets, to gain a much deeper insight into what makes markets tick and how the fundamental school interacts with the technical school.

In fact, as the years have gone past and my market experience and analysis has developed I have come to realise something that I believe many market participants do not fully appreciate. Most technical market analysts or investors started out as fundamental market analysts or investors before, having become disillusioned with the fundamental school, steering themselves in the direction of the technical approach. However, most fundamental market analysts started life as fundamental market analysts and dismiss the technical approach without ever having studied any of the theories at all! This, I think, is extremely instructive and is one reason why market price action behaves and cycles the way it does. More of that later.

So having completed my economics degree I went to work for a stockbroker where I gained my first market experiences of managing stock portfolios by analysing company accounts, profit and loss statements, attending corporate meetings and taking account of news flow. Younger readers may be shocked that our desk of three people had to share a quote screen for the stock prices we monitored! That was 1991! How times have changed. While working there I actually believed that reading the

financial press, diligently studying the company accounts, talking to the people in the company and following the news flow from the company would help me to predict the future course of the company share price.

And this is when it hit me. Why, when all the news was positive, the company was making money and the balance sheet was in healthy shape, would the share price go down? Why, when all the news was negative, the company was losing money and the balance sheet was a joke, would the share price go up? This was not how the markets were supposed to work for goodness sake. The textbooks all told me that if a , b or c happened then logically x , y or z should happen to the share price, but the only problem, it seemed, was that no one had decided to inform the markets that this was the case! I was intrigued and determined to find out as much as I could about the reason for this beguiling market behaviour.

As luck would have it my next job, in London, was managing bond and currency portfolios with Hambros Bank – one of the last of the venerable blue-blooded British merchant banks in the City – and it was here that I worked closely with a true giant in the technical analysis field, Tony Plummer. It was an honour to work with someone who took a much more rigorous intellectual approach to technical analysis than the majority of his peers and his classical work on price cycles became a major influence on my market analysis. I joined the UK Society of Technical Analysts, learned all I could about the subject and passed the diploma in technical market analysis. Now, one of the most undermentioned aspects of technical analysis is just how broad a subject it is. Essentially, technical analysis is concerned with using past price data in order to anticipate (anticipate rather than predict!) future price action but this encompasses a massive range of methods and theories. Dow Theory, Japanese Candlesticks, Elliott Wave Theory, Cycles, Gann Analysis, Behavioural Finance, Point and Figure, pattern recognition, trend lines, Chaos Theory, quantitative analysis, the hundreds and hundreds of price derivative indicators and Trend Following are only a part of the subject that is known as technical analysis. Add in sentiment indicators such as survey and positioning data and you have something that is massively broad in scope. Yet many people still think of technical market analysis as quite a narrow field. It is not. The term *technical market analysis* refers to anything that involves a study of the market price action as well as other related indicators, such as volume, to gauge the sentiment and psychology of the overall market in order to take advantage of trends and discernible patterns in price that the psychology cycle emits. The subject is extremely broad and the theories that underpin it pre-date modern theories such as the Capital Asset Pricing Model (CAPM) by centuries. As I wanted to incorporate such theories into my overall investment approach, I worked with technical analysis throughout my time in London while at the same time paying attention to the so-called fundamentals of the market.

By this stage in my career I had armed myself with a very broad understanding of all aspects of market analysis from the fundamental approach to the technical approach. I had seen throughout my career that some forms of analysis worked well at certain times but then didn't work very well in the next period, and those that

gave very bad results would all of a sudden come into their own and produce stellar performance. It was for this reason that I still struggled with the age-old question, the grand daddy of market analysis, the big kahuna of them all, that anyone who has had any interest in trading or investing in the markets has asked themselves at some point of their lives.

That question very simply is . . . *what works?* What method of analysing the markets works all, if not most of the time? What is the Holy Grail of market analysis and trading or investing? If we are honest with ourselves, anyone who has a passion for the markets would admit that this question has vexed us all in our quest for reliable and consistent profitable trading or investing and I am sure it will continue to do so because it is, at the end of the day, human nature to look for the consummate answer to puzzles and things that we, as human beings, seem powerless to control. And the markets are by far and away one of the biggest puzzles known to man. The question will also keep the steady supply of books, courses and what not with titles like *Trading Secrets* and *How to Make a Million Dollars in the Markets* coming in and being snapped up by people searching for that elusive Holy Grail.

After my stint in bond and currency market asset management I went on to concentrate full time on currency markets specifically. The currency or FX (foreign exchange) markets are sometimes referred to as the “most efficient” markets. This is confusing for the layman because the word “efficient” here is being used to describe the ease of dealing or the liquidity of the markets in terms of the bid (selling price) and offer (buying price) spread. At nearly three trillion US dollars turnover *per day* there is certainly enough liquidity to dwarf all other markets in that respect, but in terms of “efficiency” in the academic sense the currency markets are maybe the least efficient markets.

The efficient market hypothesis, the academic idea that information is absorbed into the price almost immediately and that, as a consequence, fundamental and technical analysis is useless, has been rightly condemned by practitioners in the markets. What the efficient market hypothesis is essentially saying is that prices move randomly, and therefore there should be no trends in the markets, but one only has to glance at a chart of a financial market to see that this notion is complete nonsense. Markets *do trend* up and down. Admittedly they do not trend up and down all the time and they sometimes trend sideways! This, in fact, is what this particular book is about.

Now, one of the assumptions that the efficient market hypothesis makes in coming to the conclusion that information is almost instantly digested into the price is that the market participants are all what is known as profit maximisers. That is, the market participants’ motivation for trading or investing in the market is to secure a profit. This assumption might well be true for markets like stocks and bonds, but when we examine the diverse range of participants involved in the foreign exchange markets then the assumption breaks down. The foreign exchange markets are made up of some profit maximisers like bank traders and currency investment managers but the market also has, among others, governments whose role in the market can be

to attempt to smooth out fast moves, industrial and commercial corporations whose actions in the market can involve buying and selling for import and export purposes and tourist companies whose actions are the result of retail tourism decisions. So with not everyone in the foreign exchange markets being profit maximisers, the informational importance contained within the efficient market hypothesis gets diluted a great deal and the market transactions, at any one time, can be the result of a multitude of very different decision criteria. Therefore, the FX markets are actually one of the least efficient markets in the academic sense and this partly explains why the foreign exchange markets can show a greater deal of persistence (or trend) than other markets as information seeps into the foreign exchange markets, or more accurately is acted upon in the market, contrary to the popular image of the FX markets being lightning speed and dynamic, at a much slower speed than the efficient market hypothesis would suggest.

Realising this, however, did not help me in my career-long search for the illusive Holy Grail of trading. "Foreign exchange markets trend more than other markets." Great, I thought. So employ a trend-following method and watch those profits roll in. Hhhmm . . . not quite! Yet again I found that there is not one method of trading or investing in the markets that is consistently profitable in every time period. Sometimes the best results came from following the trend, sometimes the best results came from fading the trend, and sometimes the best results came from doing absolutely nothing at all!

Therefore, from where I sit today, having managed institutional money in stocks, bonds and currency markets (and my own private money in commodity markets) I feel that I have a pretty good understanding of what makes most markets tick. After years of painstaking research, academic testing and a huge amount of real world trading of methods and theories trying to work out the best way to trade and invest in the markets, I have come to this not so startling conclusion. *Everything works . . . some of the time.*

This is what this book is about. It is about a pragmatic approach to trading and investing rather than a "trading method that will make you rich" type book. This book is not about a "guaranteed" way to make money and is certainly not about some sort of "magic" system that will make you a millionaire. If that is what you are looking for, I suggest that you stop reading now. Incidentally, there is actually a book called *The Trading Rule That will Make You Rich* by Edward Dobson and I thoroughly recommend it because it has some very valuable insights about market behaviour. However, as will become clear to the reader, I would not recommend trading all your capital on just one methodology.

The realisation that "everything works but only some of the time" is, in my opinion, crucial to long-term success or survival in the financial trading market place. Over the course of my career in the financial markets I have found out that there is actually no holy grail for trading success and there is no magic method that will guarantee you to trade or invest consistently profitably. Instead, what I have found is that by analysing the probabilities of what *trading environment* or "*regime*" the

market is in, or about to be in, can enable you to adjust your trading or investment risk parameters and game plan in such a way as to increase your chances, your probabilities, of long-term success. And probably the best barometer of success in the markets is the longevity of staying involved. To paraphrase a quote from Jesse Livermore, probably the greatest trader who ever lived, the real aim of the investment or trading game is “to be able to come back tomorrow and play it again”.

HORSES FOR COURSES

On the subject of games, Lee Trevino, in his day, was a great golfer. He won six major championships (the four top competitions in the sport) and he was known as one of the best strikers of a golf ball who ever lived. He could hit low, boring shots into a fierce wind and he could hit high floating shots that would land on the green like, as he said, “a butterfly with sore feet”. There was however one thing Lee Trevino never did. He never won, in fact never really played very well, at the Masters tournament played at Augusta National, Georgia, USA. The reason for this is that Trevino played his golf shots with a fade (a left to right spin on the ball for a right hander), whereas the layout of the golf course at Augusta dictates that golfers play a lot of draws (a right to left spin on the ball). So Trevino’s game was simply not suited to the Augusta National layout.

Pete Sampras was one of the best, some say *the* best, tennis players that ever lived. He won a record 14 Grand Slam men’s singles titles, for six consecutive years finished ranked number one on the ATP rankings and *Tennis* magazine named him the greatest player from 1965 to 2005. Yet, despite all this, there was one thing Pistol Pete never did. He never won a significant title on the clay court playing surface and the generally accepted wisdom is that this was because the slower surface didn’t suit his natural serve and volley game.

In the business world, the commodity sectors in quoted stock markets are, on the whole, filled with very well run companies. They sell a product that has a limited supply, keep costs under control for the most part and are, with a few exceptions, generally prudent in not over-extending themselves with debt. However, everyone knows that the share prices of commodity stocks will still generally underperform the rest of the market when the prices of commodities in general are going down.

The point is this. Lee Trevino was still a sublimely talented and successful golfer, Pete Sampras was still one of the all-time tennis greats, and mining companies still have successful business models over the long term. Great athletes and solid companies can be renowned as such even if they do not perform in each and every environment in which they ply their trade. Of course, when someone or something does come along that is so great that they *do* perform in each and every environment (such as a Jack Nicklaus or a Tiger Woods in golf!) then we just have to sit back and admire the exceptional brilliance of that. They are the very rare exceptions to the rule though, and if this *variable* performance behaviour can be observed in what are

essentially competitive endeavours where the environment in which the endeavour is performed is changeable, then should it not be true for other competitive endeavours such as investing in or trading the financial markets?

Investing in or trading the markets has many similarities to sports. Both have a scoring system, league tables and defined rules. Both have direct competitors and both involve a judgement of risk versus reward. These are just some of the reasons why a lot of the most passionate people involved in financial market investments or trading have sports backgrounds or are keen followers of sports. They have a natural competitive spirit and are motivated by endeavours whereby gaining an edge on your competitors can be the difference between winning and losing. They constantly strive to be better at what they do. However, just as their sporting counterparts develop their own skill sets, they will develop a *style* of trading or investing that will suit some market conditions perfectly and others not so well. The realisation we come to therefore is that, just as racehorses perform best on certain race courses or in certain weather conditions, just as golfers perform better on certain golf courses, just as tennis players perform better on certain surfaces, and just as different companies perform better under certain economic conditions, various styles of financial and commodity market trading or investment will work better under different market conditions.

As an example, take trend following as a market strategy. Trend following is perhaps the most simple and straightforward market strategy where, as the name suggests, the strategy is to follow market trends as they develop. Identifying where trends start and end is done through various means but the most popular would be by using moving averages of the price. Let us assume for simplicity that we are using this strategy on just one market. A proper and complete trend-following methodology would include many other markets because of the good diversification benefits that accrue, but let us assume for now that we are trend following on just one market. Now, when the market in question is trending then this strategy will produce good, solid, perhaps even spectacular results. However, when the market in question is not trending then this strategy will produce bad, perhaps even disastrous, results. This is just common sense and not rocket science! It is a statement of the blindingly obvious. When this market does not trend then trend following as a strategy will lose heavily, but when it does trend the strategy will win superbly. The questions therefore become (a) do you as the investor in this strategy think that a trending market will occur more often than it does not, and/or (b) how much does the strategy lose when markets are not trending and how much does the strategy win when markets are trending?

I remember trying to explain this latter point in a presentation at one time. My presentation was on trend-following strategies and, as a believer in the strategy myself (over the long term and with appropriate diversification), I was essentially trying to persuade a group of board level people why we should use them. Now, there is a certain phrase that people use in the market that is very misleading. In fact there are a few but more of that later! This particular phrase I am referring to is one that people like to trot out from time to time in a light-hearted jovial sort of way as

if they are saying that they are experienced market people and know what they are talking about. The phrase is “as long as you are right more than you are wrong then that is all that matters”. If there was ever a saying that has misled and is still misleading market professionals then this is it. It is up there at the top of the list of market sayings that sound good but, when you actually examine the details, turns out to be the biggest load of rubbish. If the people that said this sort of guff just sat down and thought for a while about what they were saying, then they would surely realise the absurdity of the statement.

There are two elements to a trade or investment in the markets. One element is the direction of the trade you want to express. That is, do you think the market will go up, down or sideways? Do you buy hoping for an up move in the market, do you sell (go short) hoping for a down move in the price or do you put on a trade that benefits from a sideways market move? The other element to a trade is how much risk you put on the trade. This element is actually the most important part of the trading or investing decision yet it is discussed very little in the sea of literature on the subject. It is this *how much* question that will have the biggest effect on your trading decisions. Let's say you are bullish of the stock market and you buy S&P Index futures. If the market goes up 10% and you have bought only one futures contract you will have a very different economic outcome than if you had bought five futures contracts. I apologise if the reader thinks that this is so simplistic that it borders on the condescending but I honestly believe that most people in the markets do not fully appreciate this money management (or bet sizing) element to trading or investing. We spend the vast majority of our time concerning ourselves with timing the market as to when to get in or when to get out when, in reality, the biggest determinant as to whether the trade or investment will be profitable or unprofitable is how much risk we put on initially, and how we manage that risk through the life of the trade. Much, much more on bet sizing later.

So, there I was presenting on trend-following strategies to these old colleagues of mine and had gone through various performance statistics over and over again to emphasise the point that trend following tends to work over the long term. One particular old boy, who had, in common with a lot of people of that generation, spent his entire career with the same company and also had a worrying habit of falling asleep in meetings, arose from a slumber and said “yes Murray, these trend-following strategies are not very good you know because only 30% of the trades are right. In this business you have to be right more than you are wrong.” Fantastic I thought. Here was a chance to really emphasise why trend following is a good long-term strategy and went on to make the point that although the system loses 70% of the time the actual amount that it loses each time it is wrong is far, far lower than the amount it wins each time it is right. The average win is much, much higher than the average loss on good trend-following strategies and that is why, over the long term it tends to be a good strategy.

However, no matter how many times I reiterated this, the man in question simply could not get the point. I think his colleagues got it but we spent many, many minutes

discussing the point and this man still repeated the mantra of “you have to be right more than you are wrong”. It really underlined the fact to me that many people in the trading or investing world have such closed minds that they simply cannot comprehend anything that goes against their perceived bias. The problem is that most people take the expression of “you have to be right more than you are wrong” to mean that you have to be right more *times* than you are wrong and this is plainly incorrect. It is a mathematical fact that you can be wrong more times than you are right, but so long as your wrong trades lose, on average, less than your right trades then, over the long run, you make money.

Nevertheless, the investment world is filled with people who have closed minds and an unwillingness to adapt, and it is this stubbornness of attitude which turns out to be good news for trend followers because by its very nature trend following forces the investor to align himself or herself with the forces of change. As Charles Darwin so famously said, “It is not the strongest of the species that survives, nor the most intelligent that survives. It is the one that is the most adaptable to change”, and so it is the people that take risks and change that take advantage of the people that do not. It is having this ability to change when, as the economist John Maynard Keynes famously said, “the facts change, which is critical to long-term success in any business, and particularly when it comes to trading or investing in the markets having the ability to accept, embrace and *follow* change is absolutely essential”.

Charles Sanford was at one time the head of Bankers Trust, a company that in its day was one of the biggest and best risk takers in the markets. He made a speech in June 1989 at the University of Georgia, USA, to college graduates in which he quite distinctly and superbly summed up the nature of risk taking in relation to what it means in the context of change. Here is an extract from that speech.

From an early age, we are all conditioned by our families, our schools, and virtually every other shaping force in our society to avoid risk. To take risks is inadvisable; to play it safe is the counsel we are accustomed both to receiving and to passing on. In the conventional wisdom, risk is asymmetrical: it has only one side, the bad side. In my experience – and all I presume to offer you today is observations drawn on my own experience, which is hardly the wisdom of the ages – in my experience, this conventional view of risk is short sighted and often simply mistaken.

My first observation is that successful people understand that risk, properly conceived, is often highly productive rather than something to avoid. They appreciate that risk is an advantage to be used rather than a pitfall to be skirted. Such people understand that taking calculated risks is quite different from being rash.

This view of risk is not only unorthodox, it is paradoxical – the first of several paradoxes which I’m going to present to you today. This one might be encapsulated as follows: playing it safe is dangerous. Far more often than you would realize, the real risk in life turns out to be the refusal to take a risk. In other words, the truly most threatening dangers usually arise when you shrink from confronting

what only appear to be the most threatening dangers. What is widely regarded as “playing it safe” turns out not to be safe at all.

I’m suggesting that you take a positive view of risk . . .

. . . We all know that modern civilization owes much to the ancient Greeks. As the 20th century draws to a close, it’s difficult to single out a Greek thinker who speaks more directly to us than Heraclitus. “All is flux, nothing stays still,” said Heraclitus some twenty-five hundred years ago. “Nothing endures but change.”

Most of us have come to believe that “nothing endures but change”, but its consequences still deserve some reflection. Obviously, if change is the fundamental rule of life, then resistance to change is folly – doomed to defeat. Just as obviously, if change is our constant, then uncertainty is an inescapable part of our lives. Uncertainty is unavoidable. Life is unpredictable. The very essence of life is the unexpected and the unintended, the unanticipated turns which we may metaphorically ascribe to Fate or Destiny or Providence.

Therefore, unless we wish to be tossed about like so much flotsam on the waves of inescapable change, we must place ourselves squarely in the midst of change. We must learn to ride the current of change rather than to swim against it – although people who haven’t taken the trouble to learn how the world really works will think we’re doing exactly the opposite.

In other words, risk is commonly thought of as going against the current, taking the hard way against high odds. In a world of constant change, however, a world where Heraclitus said we can never step into the same river twice, taking risks is accepting the flow of change and aligning ourselves with it. Remember the first paradox: risk only looks like reckless endangerment. For those who understand reality, risk is actually the safest way to cope with a changing, uncertain world.

To take a risk is indeed to plunge into circumstances we cannot absolutely control. But the fact is that the only circumstances in this life that we can absolutely control are so relatively few and so utterly trivial as hardly to be worth the effort. Besides, the absence of absolute control – which is impossible in any case – does not entail the absence of any control, or even significant control.

There, again, is the paradox: in a world of constant change, risk is actually a form of safety, because it accepts that world for what it is. Conventional safety is where the danger really lies, because it denies and resists that world.

I trust you understand that when I say risk is actually safety, I’m talking about a certain sort of risk. I’m not advising that you leap off tall buildings in the hope that the operation of constant change will reverse the law of gravity in mid-flight. I’m speaking rather of a sort of risk which actually aligns you with the direction of change.

[Reproduced by Permission of Charles Sanford]

This speech, in my opinion, is one of the most eloquent portrayals of what the trend-following investment strategy is. Charles Sanford speaks the language of the risk-taking pragmatist. The pragmatic person takes account of change, accepts it for what

it is and instead of trying to fight against it he follows it and tries to take advantage of it. Things change. Life changes. Trading and investing market conditions change. This is why trend following is such an attractive proposition not just in the investment world but also in life itself. For those willing to *surrender* themselves to the forces of change without feeling the need to understand all the reasons why that change is happening, they will, over the long term, have the wind at their back and, contrary to the popular opinion of trend following, actually be *ahead* of the majority when important, long-lasting change happens. The expression “one must move with the times” is another way of saying the same thing. To surrender yourself and accept change, even if you do not understand why it happening, is to find contentment. Fighting change and swimming against the current is mentally and physically tiring.

At this point readers might be thinking that I am very biased towards a trend-following style of investment and they would be right. Do not fear though. This book is not a sales pitch for trend following despite the fact that, in my opinion and experience, a diversified, long-term trend-following portfolio is not only the most logical way to invest but is actually, as Sanford alludes to with regard to risk and change, the safest way to invest. Long-term trend following is the only investment methodology that I know that guarantees the investor to be on the correct side of a sustained movement in the price. Think about that for a second. It *guarantees* the investor to be on the correct side of a sustained movement in price. When I first started getting involved in trend following I thought this was an amazing statement and I still do. If the key to long-term success in the markets is, as many investors no matter what their own personal style bias is would agree, to run your winners and cut your losers then trend following is the very embodiment of such a philosophy.

However, trend following, like any other investment strategy, needs to take account of change in other ways like changes in market conditions. As has been noted, certain trading and investing strategies work well in certain circumstances and not in others so the challenge to the market participant, as I see it, must be to try to detect changes in those market circumstances. Once they have been detected the participant then needs to accept the changes for what they are and embrace them in order to fully benefit from the ebb and flow of the markets. I may be biased towards a trend-following philosophy and methodology just as others might be biased towards a fundamental approach to trading or investing in the markets, but one thing I have learnt over the years is that having this ability to take account of changes in market conditions can enhance one's own personal investment style bias in such a way as to, at least, smooth out the losing periods or draw downs in equity that inevitably occur from time to time. Moreover, if no inherent style bias exists then an analysis of changing market conditions can enhance returns over the long term.

So how do we define these market conditions? Very simply, these are the general states of *volatility* under which the market in question is trading in the period under review and it is these states of volatility conditions that have generally become known as trading regimes.

TRADING REGIMES

The phrases trading regime or trading regimes can actually mean different things in the context in which it is used, so it is worth stressing the context that I am using in this book.

The *Oxford English Dictionary* defines regime as, among others: *a manner, method or system of government, the condition of a body of water with regard to rates at which water enters or leaves it and the set of conditions under which a system occurs or is maintained.*

In the financial markets reference to a trading regime has three broad contexts.

Firstly, a trading regime can be referring to the rules and regulations that are set down for something to trade under. Take a currency as an example. It can be freely traded or “floated”, it could be fixed to another currency or basket of currencies, or it could operate under a “dirty or managed float” whereby the country’s authorities attempt to keep it within a band around another currency or basket of currencies. At the time of writing, for instance, the Chinese Yuan is pegged to a basket of currencies and the Hungarian Forint has a managed float plus or minus 15% from the Hungarian exchange rate with the Euro.

Secondly, a trading regime could be referring to strategies that can be employed in the markets. In the currency markets, for example, there are many different strategies used for attempting to make money. The “carry” strategy refers to a strategy of going short a low yielding or basket of low-yielding currencies (those with low interest rates) and going long a high yielding or a basket of high-yielding currencies (those with high interest rates). The idea is that the positive interest rate differential (the carry) will produce solid long-run returns. Another strategy that would apply across all asset classes could be a value strategy where the idea is to buy undervalued assets and sell overvalued assets. This is probably the most widely used strategy in discretionary market trading. Discretionary market trading is when the decision to buy, sell or do nothing is taken mostly by human beings acting subjectively and weighing the various factors involved in the decision-making process in their heads or having the ability to overrule a more formal process. This differs from systematic market trading where the decision to buy, sell or do nothing is taken mostly by a computer and the human involvement is limited to the initial setting up of the computer model. Yet another trading strategy could be the previously mentioned trend-following strategy where the idea is to follow the trend of the market by using systematic calculations and methods such as moving averages to decide in which direction the trend is going. Each of these strategies plus a lot more could be referred to as a trading regime in certain contexts.

Thirdly, trading regimes have been referred to in quantitative and econometric modelling studies as states of macro-economic conditions such as low (high) tax economy, low (high) inflation and above (below) trend growth economy to name only a tiny amount. These “regimes” are widely referred to in studies that involve optimisation of investment portfolios using quantitative techniques such as those pioneered

analysis a subjective exercise. Of course by including hard filters of constitutional definition it can be less subjective than purely looking at the time series and saying “that’s a trend” or “that’s a range” but not, I would argue, by much. In my opinion, it is pretty obvious to the eyeballer of a time series chart whether a market is in a trending regime or a range-trading regime and, as we shall see, how one takes advantage of these regimes can differ between methodology and time frames anyway.

To be clear, therefore, a “trading regime” in this book is referring to whether a market is in a trending mode, range-trading mode or whether it is doing a bit of both. For many years my market analysis focus has been to identify when such trading regimes exist and when, more importantly, they are likely to change or switch because it is then that we, as market participants, can choose the right trading or investment techniques that go with the particular trading regime and not try to persist with something that would obviously not produce good results in that regime. The remainder of this book outlines why the markets behave the way they do, what some of the techniques and methods are to identify what type of trading regime a market is in or about to be in and how best to use these techniques to help us in our overall trading or investing game plan.

TRADING REGIME ANALYSIS FOR THE LONG TERM

Many readers may be thinking that a lot of what I have been discussing in relation to trading regime analysis is very similar to existing methodologies such as volatility breakouts, and they would be right. In that sense trading regime analysis is trying to identify exactly the same periods of low volatility or price range contraction rising to high volatility or price range expansion. On the other side of the coin, when trading regime analysis is looking to identify periods when conditions will calm down, at least in the direction of the immediately preceding trend, then this is similar to other so-called trend exhaustion techniques.

Over the past few years there has been new quantitative research in this area with regard to identifying volatility regimes and a lot of this research has been useful. However, my contention has always been that most, if not all, of the cycles in volatility can be captured by using traditional technical analysis techniques, where volatility expansion or contraction can be anticipated just as well as volatility predictive quantitative models. In fact, a lot of traditional technical analysis techniques will not only anticipate the direction of the volatility but also, in the case of an expansion in volatility, the direction of the market.

The problem, however, has been that a perception of technical market analysis both in relation to traditional technical techniques and in relation to volatility breakouts has grown whereby most people tend to think of these analyses as useful only for short-term market movements. This, as we shall see, is unjustified because whatever time frame (or fractal) is being examined, the market is being driven by the same underlying psychology that comes from human emotions. Price patterns

repeat themselves at every degree of the market. A time series chart of one-minute closing prices will look similar in certain circumstances to a chart showing one-month closing prices because the underlying driver of the market is the same human psychology that causes price trends and reversals of those trends. This notion that our more fundamentally minded colleagues have that technical market analysis is suitable only for short-term market analysis is simply flawed. Yes, the so-called fundamentals will change over the long term, but they will do so at the very best coincidentally with changes in the price of the market. More than often the market price will discount the changes in the fundamentals before they come to be noticed by market participants and so, in this regard, technical and volatility market analysis is perhaps even *more* relevant for the long term than the short term. Therefore trading regime analysis, at least as I see it, can be broadened out considerably to look at the long time frames and add value to those who follow long-term investment processes as well.

Knowing where the probabilities lie as to whether currencies for example will be in a trending or range-trading environment over the next few months will add value to a currency overlay manager who has a choice of models or styles to utilise in extracting profits from the markets. Trying to estimate the longer term probability of volatility can yield some spectacular results in terms of strong long-term trends and, in a sense, volatility analysis could be said to be most applicable to longer term analysis because strong volatility can occur at exactly the time when a market has been doing nothing for a long period of time. If a market has been going sideways in a period of very low volatility for a long, long time then it tends to get forgotten about by the market. Then what seems like all of a sudden the market springs to life and because it has been in a low volatility environment for so long the subsequent trend can be much stronger and longer lasting than the majority expect. A perfect example of this was the exchange rate of the Euro versus the Great British Pound in 2007 and 2008 when, subsequent to this period, the exchange rate had been in an extremely low volatility period for nigh on two years before moving up by over 20% in a strong trending fashion in only eight months – a huge move for a major exchange rate.

As we shall see, therefore, I have concentrated a lot on *longer term* examples of trading regime but an analysis of volatility is very applicable to all time frames.

TERMINOLOGY

Throughout this book I refer to analysts of a trading regime or volatility as, you might have guessed, trading regime analysts. I recognise that technical analysts might point out that I could just as easily use the term technical analyst for most of what I describe and I would have to agree with that. However, I would like to make it clear that what I am describing in this book is an analysis of the markets that, first and foremost, is attempting to identify and anticipate the current and likely trading *conditions* of the market, be they trending or range trading. This analysis is not as its primary goal

by Harry Markowitz. This definition of trading regime is probably the one that most economists and econometricians will think of in any discussion on the subject.

However, lastly, and this is the context in which this book is referring to trading regimes, a trading regime could be referring simply to whether the market in question is in a trending mode or in a range-trading mode. A market can do one of three things. It can trend, it can trade in a general range or it can do a little bit of both. The purpose of trading regime analysis is to identify what trading environment the market is in and, crucially, to identify when it is *probable* that the market is about to enter either a trending or range-trading regime. To me this is the essence of the markets. Markets trend, then they don't trend, then they trend again and this cycle continues on and on. As a trend follower I, of course, get frustrated when the markets do not trend, which is actually the majority of the time, but as per the earlier comment about being more right than wrong, the times that they *do* trend tends to more than compensate for the times they do not. Nevertheless, this was my main motivation for studying trading regime or *volatility* conditions in the market because if I could identify when the market was likely to enter a non-trending period then I could take appropriate action to defend my investments because I knew that a period of range trading would be bad for trend following. By the same token if the volatility conditions of the market meant that a trending regime was more likely, then maybe I could take advantage of that by increasing the risk I was applying to the market. So in this regard, trading regime (or volatility) analysis actually assumes a very important role in determining many aspects of the overall investment process.

Now, how you may ask do I propose to measure whether the regime in question is a trending regime or a range-trading regime? Well this may disappoint the more quantitative readers but my criteria for identifying what type of trading regime the market is in is for the most part the good old-fashioned, original and arguably the best computer in existence. That is the human brain and more specifically the eyes. The good old eyeball method of analysis is, in my experience, the simplest and most intuitive method of analysis for looking at price action for the purposes of deciding whether the market is trending or ranging and, in my opinion, because I am not advocating a purely systematic, mechanical method of actually trading the markets, a subjective and loose definition of what constitutes a trend and what constitutes a range should suffice.

Many people have tried to define what a trend is and many people have tried to define what a range is, but the trouble with these definitions is that they all have the one thing in common in that the parameters of definition that go into the models were, at the end of the day, thought up by a human brain. So someone's definition of a trend could be that the market has to have an upward sloping 20 period moving average and a rolling correlation coefficient relative to time of ± 0.75 but someone else's definition could be completely different. In the same vein, someone's definition of a range might be that the market is being capped within a 1% price range high to low for five days or more but, again, someone else's definition of a range might be completely different. You see, the definition of a trend or a range is in the final

attempting to anticipate *direction* but it does include traditional technical analysis methods and pattern recognition that would, as we shall see, anticipate direction as an inherent part of the analysis. So while technical analysts will be very familiar with most of what is in this book I believe the potential use of the techniques and methods that go into an analysis of the trading regime of the markets has a slightly different bent to it; for instance, in the uses of position sizing and/or risk budgeting. I would ask the reader therefore to keep in the forefront of their thoughts when I use the term *trading regime analyst* that the primary goal of this endeavour is to look for the probabilities of underlying market's *volatility conditions* and not necessarily the *direction* of the market.

In this sense it is the probability of *volatility*, in whatever direction that occurs, rather than the probability of *direction* that is our main aim. An understanding of the current and anticipated volatility conditions will help the market analyst in many ways; for example, in terms of the position sizing of investments, the allocation of risk to certain investment styles and the choice of markets in which to invest. Volatility analysis does not have to be the realm of the quantitative community only, and orthodox technical market analysis can give the analyst a good gauge as to the current and likely trading regime.

In the next chapter we shall discuss how it is the primeval driver of human beings themselves that is the primary driver of all market price action and the subsequent volatility, or lack of it, that it creates.

2

The “Nature” of Markets

“An American monkey after getting drunk on Brandy would not touch it again, and thus is much wiser than most men.”

Charles Darwin

THE HUMAN MARKET

London, November 1988, 0600 hours.

It's a cold, dark winter's morning and Jimmy the foreign exchange trader walks into the bank where he plies his trade buying and selling the world's currencies for profit. His normal routine is to get in very early before any of the other traders and brokers arrive so that he can prepare for the day ahead, catch up with what has been happening in Tokyo and read some research. Having read the newspapers and some research he spies a copy of *The Economist* magazine and picks it up. About 30 minutes later his colleagues start to drift in to the trading floor where they find Jimmy busily buying Japanese yen. Jimmy talks to them and they start buying yen too. The price of yen starts to rise. After a while, more of his colleagues arrive at work and they too start to buy some yen. The price of yen rises even more and, as the London market gets into full swing, the price of yen rises even faster. Two hours later the price of the Japanese yen has risen so far that Jimmy and his colleagues start to close out their long yen positions for many millions of dollars profit. Jimmy's boss, who has been standing at his office door next to the trading floor watching all this, comes over and asks Jimmy why he had been buying yen. “This”, he exclaims as he hands over *The Economist* magazine he had been reading, “tells me that the Japanese economy is growing very quickly, the Bank of Japan is thinking of raising interest rates and that the Japanese trade surplus is so massive that the American car manufacturers are lobbying the government over the cheapness of the Japanese currency. The yen can only go one way and that is up.” The boss looks at the magazine and then looks back at Jimmy. “I congratulate you Jimmy”, he says. “You have made this

selling prices, supply and demand. It is human beings who are market makers and have to determine how much to mark prices up or down in response to the supply and demand levels they are seeing in the market place. The dealing price we see on our screens hasn't just been made up by a computer – it represents the level at which a human being is prepared to strike a deal. In fact, the indicative dealing prices we see on our screens sometimes *are* made up by a computer, but when we come to actually deal we find that the selling price can be very far removed from this indicative level. This is because a human being has to make a *decision* and quote us a “firm” price – the price he is committed to dealing at if we say “deal done”. This reality brings it home to everyone in the market place that it is human beings who ultimately have to decide whether to accept risk or not. That is, it is human beings who ultimately have to decide on the price at which they are prepared to deal, and it is therefore human beings (and their emotions) who ultimately decide the direction of prices.

A prime example of this process of price discovery occurred in August 2007 as global stock markets were hit hard by a crisis in the credit bond and derivatives markets. The problems were brewing for a few months but came to a head during July and August when players in the market actually wanted to sell some of their positions but found out that no one wanted to buy them! There was “no bid” as they say in the markets. As a result, confidence in the actual functioning of the financial system had declined and many people who owned assets in this area had been forced to admit that they were worth a lot less than they thought they were worth. Generally the way market assets are valued is a process called “mark to market” whereby at the end of every day the assets of a fund or a bank are marked to a market price. The assumption of course is that this market price will be close to, if not exactly, the price that *would* be achieved if the assets were to be sold there and then. The problem with the credit bond and derivatives in the months leading up to August 2007 was that people were not marking to market but they were valuing the assets on what a computer-generated program told them they were worth. Well, call me old school but I have always believed that something is only “worth” what someone is prepared to pay for it, so it is madness to value something on what a computer tells you it is worth. What these people were doing was not marking to market but “marking to model” or, as someone hilariously put it, “marking to myth”!

If they had actually marked their assets' prices to a price that was a fair reflection of what someone else was prepared to pay for it (not what *they* thought someone *should* be prepared to pay for it!), a lot of the sudden downturn in stock markets could have been avoided. Of course their funds would have started suffering a few months earlier but better that, when they can still salvage the situation, than to wake up one morning when you are on holiday in the Hamptons and find that your hedge fund has gone bust in one week! In military combat, is it not better to be wounded badly with still an outside chance of survival than to be blown to bits in a second?

It is therefore the market players, not computers, who decide what the actual buying and selling prices are for assets, but I think a lot of people in the markets tend

bank a huge amount of money this morning. There is just one thing old boy,” he says as he hands the magazine back, “your copy of *The Economist* is 12 months old.”

This is an old City of London legend and I have no doubt that it is actually true! As well as being slightly amusing it is also a very instructive insight into how financial markets actually work. Let’s think about what happened here. The trader comes in, reads something that he construes to be positive for the Japanese yen and, crucially, decides to act on his thoughts and he buys yen. This very *act* causes a demand for yen to be seen in the market place by other traders or market makers. A market maker is someone who is obliged to quote, buy and sell prices to people that want to deal. He is a “price maker” whereas other traders and investors who do not have such an obligation are “price takers”. They are taking the price that others quote them. The market makers see that there is demand for yen in the market and, when Jimmy’s colleagues also decide to start buying yen they see even *more* demand. What do they do? Unless there is the same or a further supply of yen coming from somewhere else, the market maker will *have* to raise the level of his bid and offer. The market makers are seeing continued demand for yen so they raise the price of yen to protect themselves (they have to sell the yen to the buyer, in this case Jimmy and his mates, and then scramble around to cover their own now short position in yen). The price of yen therefore goes up because the demand that Jimmy’s bank has created for yen causes other participants in the market to quote higher selling prices for yen and for other buyers to join in. The only relevance that the magazine article has is that it caused Jimmy to pick up the phone and start buying yen in the first place. He could quite easily have sold yen due to a different set of “reasons” and we would have had the same outcome in reverse.

Which brings us to one of the, in my opinion, absolute truths of market dynamics and behaviour. It doesn’t matter what the underlying reasons are for putting on a trade, the only thing that *ever* matters to the direction of price in a market place is the actuality of trading and the size of the trade. Market makers and other market participants couldn’t care less if the yen was being bought for valid or invalid “reasons” and they couldn’t care less whether they agree or disagree with your reasoning. All they care about is that there is huge demand for yen and no supply! If they don’t mark the price higher they will lose money because they will have sold yen at a lower price than the one at which they themselves can buy it. It is that simple. Reasons have nothing to do with it. It is the pure laws of supply and demand in dynamic action and it is due to this fact that the price we see on our dealing screens has no emotion and no bias. It is what it is. It is reality.

This process of “price discovery” is the way any market behaves. A market has buyers, sellers and people who do nothing. If the buyers demand is bigger at any particular time than the current supply then the price will rise. This is what the economic textbooks tell us about how a market should function but what they do not tell us, and what is absolutely crucial to an understanding of market behaviour, is that it is the human psychology behind the numbers that is the main driver of price discovery. It is human beings who are interacting with each other to match buying and

to forget this. How many times have you heard someone say “the market *should* do this if that happens” or “the market *should* drop if these economic numbers come out weak”? This is all nonsense because an economic textbook does not govern the market. Human beings govern the market and, as such, it is not the economic or corporate numbers that move the market, it is how human beings *react* to the data or information that is the ultimate key.

The big mistake we all make is to assume that the market will react the way we would react – that is, in a rational (or what we think is rational) fashion, when the market as a whole can be very irrational. Let’s take an example of an economic statistical release like GDP (gross domestic product) and assume that it is the GDP of the USA that is being released to the market. The standard expectation of a rational investor is that if the US GDP comes out better than expected, this will be good news for economic growth and therefore good news for the US stock market; but if GDP comes out worse than expected, this will be bad news for economic growth and therefore bad news for the US stock market. So let’s say the GDP number comes out better than expected. You, as an individual, might say “OK this is good news let’s buy stocks”. But what if you already own stocks? In fact what if you already own a lot of stocks and own no cash with which to buy any more stocks. You cannot buy any more stocks even though you think the economic release is good news for the market. Now let’s assume that there is a lot more people in the same boat as you who are already long a lot of stocks. They, too, think that the better than expected GDP is good news for the economy and therefore should be good news for the stock market. But they, too, cannot buy any more stocks because they have no spare cash. Now what happens? The only way the stock market will go up after the release of this economic report is if there is actual demand to buy stocks. Sure the market makers might mark stocks up a little on the expectation of demand coming through, but the price will have to move back down if no demand actually appears in order for the market to “clear” prices. We then have a situation here where an economic release has come out which is good news for the economy and therefore should be good news for the stock market, but the *reality* of the situation is that very little actual buying of stocks occurs.

And what if there was someone in the market who is a relatively big mutual fund player but has been experiencing a lot of redemptions from his fund. He has to sell some of his stockholding in order to pay the redeeming holders their cash and decides to do the transactions around the time of the GDP economic release. Before you know it there is actually a lot of net supply in the stock market and the market makers are forced to mark prices down. The US stock market *falls* even though US GDP has come out better than expected and *everyone* as individuals agree that this is good news for the economy and the stock market. This is when people start to scratch their heads and look for “reasons” why the stock market reacted the way it did. Pretty soon we have a list of market commentators citing this reason or that reason for the market fall when it all simply boils down to the fact that people didn’t have the ability to buy any more stocks and someone else was selling for cash flow reasons. The *actual*

market actions of human beings, not what they or you think as individuals, are the only thing worth worrying about as this is the *real* driver of market prices.

BEHAVIOURAL ANALYSIS

“Never make predictions, especially about the future.” This quote that is, depending on who you believe, attributed to either the late and eminently quotable baseball manager Casey Stengel or the other baseball icon Yogi Berra, is extremely poignant when it comes to the financial markets and especially for trend followers such as me. The blinding wisdom in this humorous quote is, however, so much at odds with what we do in the investment world. Another one of these market sayings I mentioned earlier that appears to have ingrained itself in the psyche of the financial field, and this one portends specifically to the currency markets, occurs when people are asked about forecasting exchange rates. The stock answer from a lot of people is usually a guffaw followed by “forecasting currencies is a mug’s game”. They say this because the standard belief is that a great many factors can affect exchange rates, so trying to predict where they might go in the future is too difficult to attempt. Now, it is absolutely true to say that a great many factors can affect the movement of exchange rates (or more accurately affect the psychology of people trading foreign exchange) but is it not also true to say that a great many factors can also affect the movement of other markets? It really makes me laugh when I hear people say that forecasting currencies is a mug’s game because, on the one hand, they are saying prediction of market prices is *impossible* but on the other hand they are saying some prediction of market prices is more possible than others! It’s as if they are saying that to predict the future course of a share price is easier to do because there are limited factors that drive the price, such as the earnings outlook and the balance sheet. If the earnings are positive, the balance sheet is healthy, and if the other limited factors that appear in this utopian world are all lined up, then the share price will go up. End of story. Nothing is surer. The market price will follow this rational analysis. It simply has to.

Yeah, right. Enron anyone?

What is even better about these people who repeat the mantra of “forecasting currencies is a mug’s game” is that they use this to dismiss the idea that you should manage your currency risk. So they don’t do any hedging at all on their foreign currency exposure but miss the point entirely that, to hold a neutral position on a foreign currency investment, you should be 50% hedged on that currency exposure. Being 0% hedged on the currency exposure is just as risky as being 100% hedged on the currency exposure. This is just common sense. I have been working in the currency markets for a long time and that has encouraged some of my friends and family to ask me whether they should buy their foreign currency for their upcoming holidays (for some reason they think I will be able to “forecast” what the currency will do!). My stock answer is that I, like everyone else in the world, do not have a clue of the exact path a currency will take between now and your holidays so you should buy

half the currency now and half just before your holiday. In that way they are perfectly hedged. If the currency goes up, they will feel relieved that they bought half, but if the currency goes down, they will feel relieved that they are still going to buy half at the cheaper rate. Common sense, but as someone once said, common sense is not so common.

Look, the basic fact is that forecasting the future course of prices in financial and commodity markets is one of the most impossible and futile endeavours ever to be attempted. So why try? Now let me make myself clear here. In my opinion there is a clear distinction between *forecasting* future economic developments or market price movements and *anticipating* these future developments. It may be semantics to some but to me there is an important difference and it has everything to do with belief system and cognitive dissonance. Forecasting implies a confidence level based on a belief system that the world will unfold as you, the forecaster, has determined. It also implies a level of arrogance because of this belief system as there is very rarely any concession to the possibility that the world might *not* unfold in the way you have forecasted. Although a lot of economists and fundamentalists would claim that all they are ever doing is giving scenarios that have certain probabilities attached to them, there is very rarely a plan B that is highlighted in their “forecasts”. Most importantly, for them, their central scenario is the one they become *wedded to* emotionally, and once that happens the cognitive dissonance process really kicks in. *Anticipating* future price movements on the other hand implies a level of confidence that takes into account the possibility that future price movements might *not* turn out the way initially anticipated. The anticipator of price movements, unlike the forecaster, is saying that, “I anticipate prices to move in this particular direction but I am open to the possibility that they may not”. To me this is again, simply pragmatic common sense.

Amazingly, however, to a lot of people in the investment management business this is not pragmatic common sense but a real sign of weakness! Unfortunately the investment management business, certainly on the pension and mutual fund side, tends to be dominated by a belief that “forecasting is what we do” and “we must be strong and confident in our forecasts”. To me this is not only utter baloney but is, in fact, very dangerous. I would have thought the job of an investment manager first and foremost is to protect the value of the customer’s money, not to try to become the “best” at forecasting. Forecasting implies a level of arrogance that the world will unfold the way that the forecaster has guessed, and because of this underlying “belief” it then becomes very easy for the forecaster to become wedded to his view, and being wedded to a view when you are trading or investing in the financial markets is probably the quickest way to lose all your money. People letting losses spiral out of control have caused every financial market blow-up in history and normally the psychology behind this spiralling is that the “view” will come right. When the market starts to go against forecasters they cannot comprehend why the market would be doing what it is doing. Do they change their view when this starts to happen? Of course they don’t. Admitting that the market is going against their original forecast because they might

have missed something or their analysis was wrong is anathema to these forecasters because they have been conditioned to be *confident* in their forecasts. They have done all the groundwork on the company, so there is no way the share price can go the other way. Changing their view when the price starts to go against them would be, to them, admittance that they were wrong and admitting that you are wrong is simply beyond the wit of these forecasters. But when you become a student of market history you realise that it is the ability to admit you are wrong that is, in my opinion, one of the traits of the all time great investors and traders.

It really is amazing how the press and the industry condition investors to believe that the best performing fund managers are the best “forecasters”, as if these star fund managers have the alchemic ability to see into the future that mere mortals do not have. In actual fact most “forecaster” performances tend to swing quite wildly and this is due in large part to this underlying psychology of being wedded to one’s view. When the market is confirming their view then all goes well and their performance is very good, but when the market is not confirming their view then performance suffers a lot because of this inability to change course and admit that their view is wrong. Consequently they are either going to be really right or really wrong, which I suppose is fine if you want a lot of volatility in returns (and to be fair I agree you need to have *some* volatility in returns) but it also means that in extreme cases the downside volatility periods could be ruinous.

I once had dealings with a company that almost went out of business in large part due to this inability to admit that a view was wrong and I particularly remember a conversation I had with the chief investment officer one day. I had been making a presentation on some subject (probably trend following) and during that presentation I had apologised for something trivial like not having any handouts. After the meeting this chief investment officer came up to me and said: “Nice presentation Murray but just one thing, never ever apologise for anything.” I was completely and utterly dumbfounded by that statement because I knew that someone who held an attitude that contained that amount of sheer arrogance would almost certainly have the strong psychological trait of not having the ability to admit to a mistake. I had seen this cognitive dissonance trait in many of my fundamental colleagues throughout the years but I did not expect to see such a level of arrogance and psychological flaw in someone who held such a vitally important role in a large investment firm. In hindsight, of course, I realise that this was naivety in the extreme because, generally, the people that hold such leadership roles all tend to have high ego and denial psychology. So since that conversation took place I was nervously waiting for the time that this psychology of the chief investment officer would present itself in all its glory, and that time came in a subsequent deep bear market when, having been overweight equities relative to bonds and cash in the run up to the bear market, that “view” and position was held for far, far too long when the stock markets began their long march lower. The chief investment officer was not solely responsible for holding that view as there were various other factors and committees at work, but he was a key decision maker in that process. Eventually the solvency of the company was called into question, and

they were effectively forced by the authorities (who admittedly had tweaked the rules a little and who were subsequently blamed for the near disaster – another important aspect of the denial psychology) to sell equities and buy bonds. Funnily enough this turned out to be near the bottom of the stock market fall.

This happens time and time again. A losing view is held for too long and then, when the pain of the loss becomes too much to bear, the view or the position is changed at what turns out to be the turning point of the market. Think Nick Leeson, LTCM or any other situation where losses or severe underperformance was occurring and you will generally find that when these positions are *stopped out*, then the market turns back in the direction of the original view. The totally ironic thing about this is that it tends to reinforce the forecasters belief that his view was right in the first place!! Incredible.

In the example above, the company was forced to sell equities so it was effectively forced to stop itself out of losing positions and, because this turned out to be near the bottom of the market, the people involved still claim that their original view of being overweight equities was correct. Never mind the fact that equities had been underperforming bonds and cash severely for nearly three years (!), oh no, our view wasn't wrong it just took time for the *market* to realise that *we* were right and those nasty authorities forced us into selling equities at the bottom of the market. Someone else *always* gets the blame. It absolutely beggars belief to think that such people are running our pension funds.

I was amused to read when writing this in the midst of another bear market that the chief executive officer of the group, who was intimately involved in the previous debacle, was quoted as saying that they were “naturally bullish” about the stock markets. Of course they are “naturally bullish” as the success of their business depends on the markets going up! There is a structural bullish bias in the stock market in general and from fund managers in particular. I have to admit that having a bullish bias to life in general is obviously a desirable position to adopt, but the problem comes when people cannot dismantle that bias when analysing the markets. I have been told on countless occasions that I am a manic-depressive, “Mr doom and gloom” or perma-bear, but the reality is that I tend to be very happy outside of the markets! Perhaps my cynically realistic approach stems from the fact that I spent the majority of my early years in the bond markets where bad news on the economy is generally good news for bond markets. Perhaps. All I know is that, in my experience of trading and investing in the markets, it pays to be emotionally neutral and that, to me, means not having a structural bias either way.

When is a view wrong? Hold a view long enough and it will eventually come right. Ah yes, another one of the market's gems of wisdom. I have actually witnessed people saying that, when their funds or positions had gone through a long period of severe losses and usually after they had been forced to stop out because the pain was too much, “it wasn't the view that was wrong, it was the *timing* of the view that was wrong”. *Excuse me?* So let me get this right. You had this view about future market direction and it didn't turn out to be correct for a long time but eventually the market

did go in the desired direction and so now you think you are a genius? It really is unbelievable that the investment management industry is full of people who actually believe this. A classic example of this type of thinking was the internet or dot com stock bubble in 1998 and 1999. Quite a lot of mutual and pension fund managers did not get involved in that ludicrous run up in stocks because they had the belief (their forecast or view) that these stocks were overvalued and not fundamentally sound companies. Of course the bubble eventually burst and it was very amusing to see all these managers come out and say “I told you so”.

Great. You told us so. Is that your job? Yes, the value of the investments were protected by the fact that you didn’t buy tech stocks near the highs and hold on to them when they crashed, but you have also just missed a huge opportunity of capital gain by not being long the tech stocks from the start of the bubble to its bursting. Should we as investors not be expecting our investment managers to take advantage of *all* the big opportunities that present themselves in the market and not miss out on some of the biggest opportunities in history due to the fact that their forecast or view did not allow it?

A famous casualty of this phenomenon was the late Tony Dye who was a fund manager with a firm called Phillips & Drew. Dye had not invested into the technology bull run and the firm was losing clients as a result. The pressure on Dye increased a lot but he stuck to his convictions that the bubble would burst and it would all end in tears, and of course the irony is that he ended up leaving the firm just weeks before the ultimate top. Dye refused to accept the fact that, although the technology and “dot.com” stocks were obviously overvalued, the market as a whole was creating one of the all time record bull markets – and his attitude was symptomatic of the fundamentalist value investor. The analogy he used to explain his decision was one of being asked to board a train that you were convinced was going to crash at some stage in its ten-station journey. The optimal strategy would be to stay on board for around five stops but if your main concern was safety then you should not board at all. The problem with this analogy is: How does anyone know how many stops the train will have in the first place? It could be ten but it could also be fifty. Unless you are actually on board the train you will be missing out on a journey that could turn out to be one of the best ever. And, as we shall see later, crashes rarely come out of the blue – there are always jumping off points before the big hit takes place. So yes, in my opinion we *should* expect our fund managers to get on that train and ride it as long as they can before it crashes. That, in my opinion, is a very large part of what they are paid to do.

Refusal to admit mistakes is not just a market phenomenon but it does seem to be correlated with situations where the person involved is, either in reality or perception, in a higher than average stress-filled environment. For instance, we see it a lot in politics where the decisions that are taken can quite literally mean life or death for some. As I write this in 2008, the decision to invade Iraq in 2003 with US and UK led coalition forces has been defended ad infinitum by the US President George Bush and the ex UK Prime Minister Tony Blair, even though a huge amount of evidence

suggests that, firstly, the reasons for the invasion were suspect and, secondly, the situation in Iraq since the invasion has been utterly miserable for most people there. Yet the longer it has gone on the more entrenched the view has become from these people that it was “the right thing to do”. The phrase “conviction based politics” sounds good but what it means in practice is that people’s convictions can lead down the wrong path and it is not until you are a long way down that path that the realisation dawns of a wrong turn having been taken a long time ago.

There does seem to be a correlation between an inability to admit mistakes and roles where decisions are transparent. It is much easier to admit a mistake when only a few people know about your decision in the first place, but when your decisions have a higher profile, as in the case of politics or to a lesser extent managing customers’ money, admitting a mistake becomes a source of acute embarrassment and a perceived sign of weakness. Personally I tend to agree with the quote from John C. Maxwell the American author and renowned “leadership expert” that “a man must be big enough to admit his mistakes, smart enough to profit from them and strong enough to correct them”. Or as Jim Fiebeg says “it takes a big man to admit when he’s wrong, and an even bigger one to keep his mouth shut when he’s right”.

COGNITIVE DISSONANCE

Leon Festinger of Stanford University published the psychological theory of cognitive dissonance in 1957. The theory states that human beings’ beliefs, knowledge or attitudes (their cognitions) will from time to time be in conflict (dissonance) with each other and that, because people do not like to have conflicting thoughts, the human brain will attempt to get rid of the dissonance. Festinger’s original observation centred around a UFO cult group who believed that the earth would be destroyed after a message of such was sent by an alien to a group member in 1956 (yes, those were the prime UFO conspiracy theory days in the US!). When the earth failed to be destroyed the group then chose to proclaim a new prophecy that the world had, in fact, been *saved* by the aliens for their sake and the group increased their efforts to convert people to the cause. The group, when faced with the reality that the world had not ended, chose to lessen the dissonance between that and what they wanted to believe by embracing a new cognition of the world having been saved by the very aliens they originally said were going to destroy it. In a very similar example a cult group of Russians who had hidden in a cave for months in expectations of the world ending emerged from it in April 2008 after melting snow had damaged the hillside hideout. Their cognition that the world was about to end was replaced with the cognition that the cave damage was “a sign from God” that they should emerge.

We humans have an infinite number of cognitions in our brain. These are our beliefs and attitudes, our views if you like. “I like to have the occasional cigarette” might be one and “I like to drive fast” might be another. Now, if we have other cognitions in our head like “I know smoking is bad for me” and “driving too fast is

against the law and is dangerous”, then what we have is cognitions (thoughts) that are in dissonance (conflict) with each other. If we had other cognitions like “I like to play golf” and “I like being outdoors” then these cognitions are considered to be related to each other and are “consonant”. We human beings have no problem at all with consonant cognitions and, in fact, many people argue that the desire to have consonant cognitions is one of the most basic of all human needs as much in fact as food and shelter.

This idea of humans being uncomfortable with thoughts that are in conflict is as old as life itself, and in 1620 it was nicely summarised by the English statesman and philosopher Francis Bacon in his *Novum Organum* work.

The human understanding when it has once adopted an opinion . . . draws all things else to support and agree with it . . . and though there be a greater number and weight of instances to be found on the other side, yet these it either neglects or despises . . . in order that by this great and pernicious predetermination the authority of its former conclusions may remain inviolate.

What Bacon is saying is that once we have adopted an opinion then we hold that opinion to be “inviolable” and therefore neglect or despise anything that goes against it. Therefore it is this basic burning desire to have consonant cognitions that means we humans cannot cope with cognitions that are in dissonance, and so what the human brain urges us to do is to get rid of or eliminate any conflicting thoughts (cognitive dissonance) we may have. There are a number of ways this is done.

One way is to ignore it altogether, so by pretending that smoking cigarettes is not bad for my health I can enjoy the occasional smoke with “a clear conscience”. If someone is presented with information that is in conflict with their view or what they already know then the easiest way to deal with it is to ignore it altogether, refuse to accept it or simply avoid that kind of information in general. Another way to eliminate the dissonance is to alter the importance of certain cognitions (thoughts). We could decide that driving fast is just so good (that we can’t do without it) or that getting a speeding ticket isn’t that important (I can pay it) and by doing this we lessen the problem of dissonance. If we alter the importance of our conflicting thoughts then our mind has less difficulty dealing with the conflict. So I can drive my car fast and not feel bad about it. A third way for human beings to get rid of cognitive dissonance is to add or create new cognitions in such a way as to overwhelm the existing cognitions. So cigarettes may be bad for my health but if “I go to the gym four times a week” then combined with the fact that “I like to have the occasional cigarette” I now have multiple cognitions that overwhelm the one that tells me cigarettes are bad for me.

To anyone involved in the markets, all of this must sound familiar. The theory of cognitive dissonance states that, as human beings, we are wired up to reject thoughts and beliefs that conflict with each other and we feel comfortable with thoughts and beliefs that agree with each other. This means that it is, for instance, comfortable

for us to look for pieces of market research that confirm our view. If I am bearish of the stock market then I will feel much more comfortable reading research that confirms my view than reading research that takes the opposite view. In fact what I will probably do is not read the bullish research in the first place and try to ignore it totally. Sound familiar? This is how we get wedded to market views. A view or a forecast about future market direction are cognitions, so anything that disagrees with those views will be in conflict or dissonance with them and, as the cognitive dissonance process kicks in, we will constantly be trying to justify them. We will constantly be trying to reject information that conflicts with our original view or look for other supporting evidence for our view in order to justify it.

The funny thing is that the idea of cognitive dissonance is well known in the markets although it is not referred to in the language of the psychology lab. We have all heard the expression “talking your book” when someone is expressing their views about the market, and what this means is that the person in question is expressing that particular view because he or she has an investment or trading position that would benefit if that view came to fruition. What is happening is that the person has one cognition (e.g. “I am long the stock market”) and so, in order for there to be consonant cognitions and not cognitive dissonance, the person will think and express bullish things about the stock market. If this person is presented with information that goes against their opinion, then they will ignore it, alter its importance (e.g. “it’s a fair point but corporate earnings are more important for the stock market”) or overwhelm it with new cognitions (e.g. “ok, but the merger and acquisition pipeline is building up”). It’s important to realise that we all have this cognitive dissonance process but if we are conscience of it, which many people are, then we can stop ourselves when we realise what we are doing and balance it up by maintaining a more open mind.

Another great example of cognitive dissonance in action in the markets is this. Let’s say an investment manager has taken a long position in a stock because he expects it to rise over the next few weeks on the back of good earnings reports. All goes well and the share price rises a little bit but then, when the earnings reports come out, they are worse than expected, the share price falls a long way down. Does our investment manager sell the stock and cut his losses? Of course he doesn’t. What he does is look for other supporting evidence that the stock will go up despite the fact that his original decision to go long the stock was on the back of the expectation of good earnings reports that have, in reality, turned out to be bad. He looks around for other bullish factors and, if he can’t find any, he will either make them up or lessen the importance of the fact (cognition) that the earnings reports have turned out to be bad. He is in cognitive dissonance because his original cognition (the earnings reports will be good for the stock) has been met with another conflicting cognition (the earnings reports were bad for the stock). So by looking for other cognitions to overwhelm this conflicting one he might say “well, these earnings reports were bad but when you look at the other fundamentals of the company they are still good and the share price should do well over the longer term”. Never mind the fact that he entered the trade as a short-term position because he expected good news. Oh no,

he now believes that the company is, in fact, a very good investment for the longer term. The share price is well below what he paid for it but he cannot bring himself to believe or accept that his original view was wrong and so he looks around for other evidence to support his long position. This is, in the words of the market, when a “short-term trade turns into a long-term position”. Again, this phrase is heard a lot in the market so it is not as if people are unaware of their behavioural traits but, despite this, a lot of people continue to submit themselves to the mercy of the cognitive dissonance process.

Cognitive dissonance is closely linked in psychology theory with denial. Denial, as the saying goes “is not just a river in Africa” but is a primary human defence mechanism and, when you think of it in this way, then it makes sense that if something challenges our views or opinions then one of the basic human instincts of defence is to deny it or ignore it. It was Anna Freud who led the way in the research of denial, and she classified denial as a mechanism of the immature mind because it conflicts with the ability to cope with and learn from reality. In the above example the reality is that the earnings reports were bad and the share price has gone down, which is in conflict with what our fund manager believed. There are various types of denial (denial of fact, responsibility, impact and even denial of denial among others) but it is denial of responsibility that is probably most closely related to the denial behaviour we see in the markets. Denial of responsibility involves blaming (shifting responsibility) and justifying (making a choice appear right despite evidence to the contrary).

The blame game has a long history in life and especially in the markets when things start to go wrong, and a classic example is, again, what happened after the internet stocks crash in the early 2000s. Investors who bought shares in these technology companies (the famous “dot.coms”) in those heady days of 1999 and early 2000 did so under their own free will. No one forced them to buy the shares (we all have our free will in the markets at all times) but they were caught up in the euphoric bubble of the time and bought at what turned out to be the wrong time (timing really is everything). The Nasdaq crashed and with it the share prices of all these dot.coms went into freefall, in some cases falling 100%! However, rather than accept responsibility for buying the shares at the wrong time (and more importantly not getting out when they were collapsing!) there began a high profile campaign to “seek justice” for these poor shareholding souls who had lost big. As a result, the investment banks that arranged the IPOs (initial public offerings) for these companies were blamed for selling them to the public in the first place with some internal e-mails being used in evidence that purported to call the shares junk and worse. Compensation was of course sought and paid, Wall Street was derided and criminal charges were brought against individuals. This happens every time there is a crash by the way. Someone has to get the blame.

Here’s a tip from me to anyone thinking of buying shares. Wall Street and all the other financial centres around the world *exist* to sell people stocks and shares. Selling is their job and their products are stocks and shares. The investment banks that launch IPOs have a vested interest to sell them not because they think it is good investment

or company but because they get fees for doing so. In fact, you could argue in a lot of cases where the company in question that is selling their stock (doing an IPO) has been very successful that the owners of that company want to “cash in their chips” and sell the company to the public. An almost perfect example of this was seen in 2007 when the private equity giant Blackstone decided to sell their shares via an IPO. Private equity had enjoyed a very lucrative number of years in the run up to 2007 as easy liquidity conditions around the world had contributed a great deal to the financing of venture capital. Billions upon billions of dollars flowed into private equity funds that then financed buyouts, mergers and acquisitions with this cash and easy credit. It was boom time all round for the private equity and corporate law firms. So what happens then? All parties must come to an end sooner or later, and it’s the people that leave at the right time, not too soon or too late, that have the best of times. So Blackstone decided to IPO their stock when the lust from the public for private equity was at its zenith.* A shrewd business move without a doubt and one that is made time and time again by principal owners of companies, but no one really bothers to examine why the owners are selling out. Of course they will say that it’s never been a better time to invest in this company or this industry. Why? Because the said company, or the said industry, has just experienced a very nice period of growth and bumper profits. Experienced! Past tense. The public always get sucked in by the marketing and the hype without asking themselves why, if everything is so bullish for the company or industry, the owners do not just keep their stakes the same without diluting them. Are the owners so benevolent that they want the public to join in the fun of making millions with them? What do you think? The owners of businesses who sell out after a particularly good period are, of course, wanting to cash in their chips and get out before the inevitable downturn in the cycle. So they go to Wall Street, do an IPO, the owners cash in their stakes in the business, the investment banks get their chunky fees and the public who buy the stock get screwed as the cycle turns. This happens time and time again yet we, the public, still get sucked into believing a story when the reality is staring us in the face. The owners are selling us their stakes because they want out before things turn sour. In the case of Sam Zell the property tycoon who sold his company in 2006, he actually took the step of explicitly telling us why he sold out in a Christmas message sung to the tune of “Raindrops Keep Falling On My Head”:

Capital is raining on my head; everything is liquid, we’re awash with cash to spend; The flood has drowned returns; ’cause assets keep liquefying, monetising, falling;

So I did me some Econ 101; the world is monetising faster every day; seems like we’ve gotten out of equilibrium; liquidity abounds, but relative yields keep falling as more capital keeps falling on my head;

* Blackstone shares were down 50% from their IPO price nine months after the launch.

We'll need income for our pension funds, where's it coming from; The world is monetising faster every day; the yields we seek won't fuel no party;

Tho' capital is raining on my head, illiquid assets alchemised; And there's one thing I know, To get back to normal, It's a long haul; that's global, yields won't improve 'till growth soaks up this liquid freefall;

Capital keeps raining on my head, so much is out there that the world is out of whack, when will we see balance back?

It's gonna be a long time 'till returns meet expectations / We need to be prepared for slim annuities . . .

Zell expected a downturn in the property market and made no bones about shouting the reasons why he was selling his stake at what he thought was the top of the market. And yet still someone bought it! Therefore, it would appear that the ancient Latin maxim of “caveat emptor” (buyer beware) is extremely applicable and relevant when it comes to buying anything that Wall Street wants to sell you!

We hear denial of responsibility all the time when things go wrong in the markets and in 2007 I heard another classic in relation to the “bank run” on Northern Rock in the UK. The Northern Rock bank had chosen to fund itself in such a way as to be very vulnerable and exposed to one particular part of the money markets and when, thanks to a number of events, that area of the money market effectively closed down the bank had to go to the Bank of England, the central bank lender of last resort, for emergency funding. This started a panic among investors, the shares collapsed and we had the rare sight of a physical bank run where people did, in fact, run to the bank to withdraw their funds! How could this happen?, the press were saying. Where were all the regulators? Why did nobody do anything?

I laughed out loud when people from the bank and the Chancellor of the Exchequer (the finance minister) of the UK started blaming the international money markets for the problems. The bank didn't get in this position because it chose to fund itself in a very, very risky fashion; oh no, it was those pesky international money markets “what done it”. A few months later, after the bank had to have been effectively nationalised, even the Prime Minister was blaming “the problem in the American markets” for the woes in the UK economy. This is classic denial of responsibility. As the PM at the time was Gordon Brown, a Scot, it made me think of the old Scots juvenile phrase that is used to highlight this denial of responsibility, which goes along the lines of “a big boy done (sic) it and ran away”. In this case, according to the “great Broon” a big American boy done it and ran away. You couldn't make it up.

It is the justifying aspect of denial of responsibility that probably lies closest to cognitive dissonance though. Rather than blame someone or something when things start to go wrong, we in the markets mostly seek to justify our position by looking for other supporting evidence. This happens time and time again in the markets, and in extreme cases such as Nick Leeson at Barings and the LTCM debacle it can have extremely serious consequences.

EGO

All of our behaviour of course is related to ego, which is one of the three divisions of the psyche in psychoanalytic theory, and the problem with ego is that it is a double-edged sword. Most of us would think of someone with high ego as being arrogant and “egotistical” and while this would probably be right it means that our perception of ego is skewed towards negative connotations. After all, arrogant and egotistical people are generally not very nice human beings. However, the human animal actually requires a certain amount of ego in order to be able to function. In order to have the confidence to do everyday tasks that most of us would normally take for granted, we need to have a certain amount of ego or, put another way, self-confidence. Getting behind the steering wheel of a car and driving to the shops requires a certain degree of ego or self-confidence (and of course a driving licence!); to apply for that job or promotion requires us to have the confidence in our abilities; and writing this book requires a certain amount of ego. This is our ego functioning and so without a certain amount of ego in all of us nothing would ever get done! We all *have* ego but the amount of ego we have or the amount of our ego awareness varies a lot. Someone who is aware of his ego will come across as much more reserved than someone who is unaware of his ego. He might be more risk averse, and while this is a good thing in times when risk needs to be controlled it can be detrimental when the situation calls for bold, risk-taking decisions. Sometimes it pays to have high ego and take a big risk, but the problem is that differentiating those times is for most people extremely difficult to do. In other words, someone with high ego will tend to take big risks *all* of the time and this becomes a problem when events are turning against him. Individuals with high ego will generally be the type to “double up” or “buy the dip” when the markets turn down because their self-confidence of being “right when the market is wrong” is overpowering. Sometimes they will be right and be heroes but the times when they are wrong is when their, and their clients, funds blow up.

WAS KEYNES A TECHNICAL ANALYST?

John Maynard Keynes. The name is synonymous with twentieth-century economics and the famous English economist is considered to be the founding father of the modern, state interventionist economics that, depending on your opinion, has either transformed the western world into one that should never again have to deal with a great depression like the one seen in the 1930s or has been responsible for creating a nanny state dependency culture where pure wealth creation is crowded out by the government. Keynes’s influence on modern economics is profound to say the least, and his writings have become absolutely required reading for all students of the subject. There are some legendary quotes from Keynes such as “in the long run we are all dead” (*A Tract On Monetary Reform*, 1923) but one quote in particular has a certain resonance with me. He was once asked a critical question about why he had

changed his view on something and he replied, “When the facts change, I change my mind. What do you do Sir?” I find it truly ironic that such a beautifully appropriate statement about analysis came from the mouth of one of the giants of fundamental analysis when it could quite simply be the epitaph on which technical market analysis stands. Not only that, but the process that Keynes describes is one that, in my experience, very few fundamental analysts find easy to live by.

For a technical analyst the facts are easy to identify. The price of the market is the fact. There are other derivative indicators of the price that are used in the course of technical analysis but the price remains the simple, basic, objective fact. When the price changes, the facts change.

For the fundamental analyst the facts are much more difficult to ascertain because there is a huge amount of subjectivity involved in coming to one’s forecast or view. Is it a “fact” that, for example, a currency is overvalued when there are so many different ways of valuing a currency from purchasing power parity to real effective exchange rates? Is it a fact that, for example, the Federal Reserve will cut interest rates at the next meeting when nobody but the members of the committee have a clue what they will do? Is it a fact that, for example, a company’s future earnings will show the projected rate of growth that the company or some analysts forecast? No, these are not facts. Mostly the fundamental analyst deals with opinions – opinions that may be based on certain measures, but opinions nonetheless. The market price is a fact.

So when does the fundamental analyst know that the facts have changed? The technical analyst knows the facts have changed because the price has changed, but the fundamental analyst has no real objective measure on which to rely. She can look at historical measures of economic statistics and corporate results but these are what the name says, history. When these facts change it is inevitably too late because the market has already moved and discounted these changes in the fundamental facts. This is the job of the market – to discount the future historical facts via the collective opinions of traders or investors today and these collective opinions are captured in one number: the market price.

A market cycles through optimism to pessimism, optimism to pessimism and one of the constants in all of this is that the peak in optimism coincides with the market’s high price and the peak in pessimism coincides with the market’s low price. Put another way, all the good news peaks at the top price and all the bad news peaks at the bottom price. The sky is always clearest at midday and the darkest hour is before dawn. So why oh why would anyone want to follow religiously something that will mostly lag the market price and produce the best or worst analysis right at the very top or the very bottom?

Another quote from Keynes’s work also alludes to the fact that he was a closet technical analyst. He said, “Successful investing is anticipating the anticipation of others”. Wow! If Keynes was not a believer in price action or technical analysis he certainly had a gift for expressing the beliefs of those analysts succinctly. Perhaps he came to his more realistic views about how markets work after proclaiming in 1927

“We will not have any more crashes in our time”, two years before the Great Crash in Wall Street (it is amazing what a little humility can do)! As has been mentioned, the goal of the market analyst should surely be to work out what the market as a whole will do and that means approaching the analysis not from the point of view of one’s own self but from the point of view of the market in general. What I mean is that analysts should be working out who is in control of the market, where does the consensus lie and either go with that flow or go against it. However they should only go against the consensus if that consensus is gauged to be extreme (i.e. if everyone is long of the market) otherwise they will be trying to sail against the wind.

Keynes’s “beauty contest” is a concept he introduced in his 1936 work, *General Theory of Employment, Interest and Money* and he used it in the context of explaining the fluctuations in stock markets. He made the analogy of “rational agents” in the stock market acting like entrants to a contest run by a London newspaper at the time whereby the readers had to choose a set of six faces from one hundred photographs of women that were the most beautiful. Keynes’s contention was that, although the natural urge of most people would be to choose the six faces that they themselves found the most beautiful, a better strategy would be to choose the six faces that the entrant thought most *other* entrants would choose, based on the knowledge of public perceptions regarding beauty. Moreover, he thought that this process could extend.

It is not a case of choosing those (faces) which, to the best of one’s judgment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practise the fourth, fifth and higher degrees.

Keynes, *General Theory of Employment, Interest and Money*, 1936

Keynes believed this process was prevalent in the stock market where investors didn’t necessarily invest on the basis of what they themselves thought about a company, but more on the basis of what they thought the market as a whole would think of that company. In this sense, Keynes described the process of technical market analysis quite succinctly. In many ways, this process and that of technical market analysis, because it aims to find out who is in control of the market (or the game) before coming to conclusions about how to act, is similar to ideas encapsulated in *Game Theory* such as Nash Equilibrium.

The technical analyst works out who is in control of the market (the bulls or the bears) and then positions himself according to which side is stronger. Surely this is a shrewd strategy to employ. I used to live in a beautiful little medieval village in Scotland called Coldingham. This was very close to the border with England and when I read about the history of the village, I learned that during the medieval wars between Scotland and England, many homes in the village would have *two* flags in their possession, one Scottish and one English, because the village would change hands so often. Depending on which army was advancing on the village, that flag

would be raised by the householders in order to avoid having their home burnt or worse. Is this the strategy of a coward? Very possibly, but it was a strategy that was as effective then in wartime as it is all over the world now in politics. Businesses and corporations all over the world buddy up with governments and prospective governments because they know that the power over a huge amount of legislation and business deals come from them. It is well known that one of the most successful media moguls in the world, who happens to be Australian, makes it his business to side with and declare support for the prospective winners in any political contest for government. The goal is always to be on the winning side, no matter what your principles might be. It's strictly business mate.

To side with the eventual winner in a conflict is undoubtedly one way to succeed and this is, in essence, what technical analysis aims to do. In the battle between the bulls and the bears in the markets the technical analyst remains neutral until she sees which side has the upper hand, at which point she declares for that side. This is just plain common sense. Fundamental analysts and economists, in general, see things through the prism of their own bias whereas technical analysts, because they are sitting on the sidelines weighing up who is going to win before acting, can be said to be unbiased. Fundamental analysts want to impose *their* view on the market and position themselves in the direction of *their own* view, whereas technical analysts work out what the aggregate winning view in the market actually is and then position themselves in that direction.

CANNONS AND TRUMPETS

Imagine, if you will, this scene.

London. 1856. The Reform Club. Lord Rothschild was sitting in the leather chair situated next to the burning embers of a crackling open fire sipping on a large glass of claret. His colleague in the opposite chair, the Duke of Westminster no less, put down his own glass on the drinks table to the side and took a deep sigh. “How in the name of God do you do it man?” said Westminster to Rothschild. “The stock market has utterly collapsed. It has collapsed even though England’s glorious military victory in the Crimea was confirmed only a few days ago and you, you (!) managed to sell out all your stock right at the very top of the blasted market. Again! That’s the third time you have done this in the last few years old boy. What on earth is your secret?” Rothschild took another sip of his claret, paused and smiled at Westminster. “It’s really very simple old boy” he said, “buy on the sound of cannon and sell on the sound of trumpets.”

That fictional conversation is based on a market legend that Lord Rothschild, the famous speculator and banker, is assumed to have described his market strategy in this manner. Buy on the sound of war cannons and sell on the sound of victory trumpets. It is often quoted in describing market action in times of conflict, geopolitical tension and war of course, but it has a much deeper meaning. The message of

Rothschild's remark is applicable to peace times too! What old Rothschild is actually saying is "buy when all the news is bad" (in his case when war is looming and/or starting) and "sell when all the news is good" (in his case when the war is over and a victory is secured). Dennis Gartman has a modern version of this phrase: "When you are yelling you should be selling and when you are crying you should be buying."

Although there is undoubtedly some truth in the fact that wars (sadly) tend to be very good business for certain sectors of industry, Rothschild's edge in investing in the markets was by sensing the psychology of the crowd and that the crowd would be influenced positively when all the news was good and negatively when all the news was bad. He realised that the markets are forward looking and driven by human perception and, moreover, that this perception is driven exclusively by human psychology. If war is looming, the news is bad, and most people sell. They are nervous about the future. As the war develops and, hopefully for the "home" side (the side where the particular stock market is), battles are being won, the news starts to turn more positive. Some (a minority) brave investors start to buy and the market starts to drift higher. Finally, when the war is over and victory is assured, the news is euphoric and the whole country is happy. This is when the last buyers come into the market – those buyers not brave enough to have bought when the news was still not entirely positive. These are the buyers who needed confirmation that the war was over and the "uncertainty" in their minds was out of the way. These buyers come in because of the perception of good news. What they don't realise is that they are the last buyers! They buy Rothschild's stock from him and, because no one is left to buy, the market falls.

This market action is also described as the "greater fool theory". Markets will only continue to go higher as long as there exists a "greater fool" who will want to buy at even higher prices. Again, it gets down to basic supply and demand. If no more is demanded at higher prices the market has to fall or at least stop going up. No matter how good the news is! *NO MATTER HOW GOOD THE NEWS IS!* If everyone is long the market, who else is going to drive the price up by demanding more?

This is the basis upon which all markets work. All the time. Over all time frames. In any country. In any economic cycle. In any political environment. It is the very essence of a market.

The truly ironic thing to me is that macro-economic and fundamental analysis in general is rooted in the theory of supply and demand yet the practitioners of these forms of analysis tend to forget that the timeless laws of supply and demand also apply to the financial and commodity markets. They look at how profitable a business is but forget that, even when the business seems to be very profitable, if there is no demand for the company's shares, the shares are not going to go up. This is the crux of the problem. They make a big assumption that a profitable company will mean demand for its shares. Of course, this link will be correct some of the time but the fact remains that they are still analysing something that is not directly related to whether the share price will go up or not. Fundamentalists of course will retort that the link between a profitable company and its share price is strong. Is it?

What about the tech bubble? Shares still rocketed, providing a great money-making opportunity, even though the companies weren’t making any money at all. They might say that currencies with large balance of payments deficits should have weak currencies. Should they? Currency markets would beg to differ on many extended periods.

In my opinion, the only true way to analyse whether there will be demand or supply of a financial instrument (like the stock market) or a commodity (like oil) is not to analyse something that may or may not be related to that supply and demand but to analyse the *actual* supply and demand of the instrument itself. If I wanted to find out about whether a football team was going to be successful would I look at their training facilities and management team or would I analyse how they are actually performing? If I wanted to invest my money with a fund manager would I base my decision solely on the firm’s investment process or would I be interested in the track record? The pragmatic answer to these questions is “you look at both”, and people generally do. Nevertheless, you can have the football team with the best training facilities and management in the world, but if they are losing games then this is reality. You can select, to manage your money, the fund manager who has the best investment process in the world, but if he is underperforming the market then this is reality. Technical analysis deals with hard reality – the *actual* price movement of the market in question. Fundamental analysis deals with soft fiction – the *assumed* price movement of the market in question.

It is this cause and effect aspect to fundamental analysis that, in my opinion, makes the markets seem so complicated and difficult to the layman. The need to explain price moves with fundamental reasons is what causes most people to get long at the top and short at the bottom. This is, in fact, why the markets swing around in their cycles and is why technical analysis will always be able to anticipate better than fundamental analysis. The need to explain price moves affects the human psyche and develops into the herd mentality. And the herd mentality is needed for the markets to act the way they do. The herd must go over the cliff at the top and the bottom – the herd act as the greater fools. The smart money has to sell to the dumb money at the top of the market in order for the market trends and cycles to occur. It is therefore the so-called fundamentals – the reasons given for why as an example a share price has moved up – that directly contribute to the herd mentality. As we shall see, the correlation between fundamental analysis and the price of the instrument is, at very best, merely coincident.

THE PRICE IS NOT ALWAYS RIGHT

One of the conclusions that the analyst must inevitably reach if he or she accepts the fact that the market price is the only fact of any real relevance because it represents the collective discounting of the future, is that the market price must always, in some sense, be “right”. For example, if the market price is going down then the

analyst must accept that the collective psychology of the market is bearish and if the market price is going up then the analyst must accept that the collective psychology of the market is bullish. In this way the analyst does not impose his or her will, or view, on the market but lets the market tell its own story and all the analyst has to do is conclude which way the balance of power lies – with the bulls or with the bears.

Over the years I have had this debate with countless colleagues and it is generally in the context of discussing another well-worn market saying – this one being “the market is never wrong”. This particular phrase or ideology actually emanated from the technical analysis school and it is easy to see why: because the technical analysis school believe that the market price contains all that is known or discounted about the future at that particular moment of time. At the same time it is easy to see why the fundamental analysis school believe this phrase is anathema because, to them, the market (price) will always present opportunities where it is out of line with the fundamentals before it moves back in line to reflect those fundamentals. My opinion on this matter is that the market price is always wrong and it is also always right! The market price is always wrong because it is always on its way to another level and is in a constant state of flux and change. However, the market price is also always right because it reflects the fact (the *FACT*) of, at that moment in time, the balance of power between the bulls or the bears. And it is the ability to always remain cognisant of this *fact* that is, in my opinion, the most important element to surviving and thriving in the trading and investment game. So the way I like to phrase it is that the market (price) might not always be right but it is always, always a fact.

RATIONALITY IS NOT A LAW

The fundamental analyst has a real problem when it comes to understanding market crowd behaviour. In general, he or she will forecast the market reaction to economic or corporate “facts” based on rational thought processes; for example, if the company being analysed announces a better than expected increase in profits or earnings, then a rational reaction would be to expect the share price to go up. This is based on what the analyst would do as an individual, i.e. buy shares in that stock. However, the big problem with this thought process is that he or she forgets or doesn’t bother to take into account that the collective market decision (as measured by the market price) may or may not assume the same rational decision-making process.

The theory of rationality is quite a grey area and tends to provoke lively debate in both academic and investment communities. Economics is a social science and is thus intertwined with humanity subjects like philosophy – and in philosophy theory, rationality and reason are the key methods used in the analysis of the data that we, as human beings, are presented with in our everyday observations. A rational decision is one where the individual or organisation concerned is said to act in pursuit of its goals and so rationality refers to the success of goal attainment.

In economics and finance the term *rationality* is used to assume that an economic agent will always make the choice that maximises his or her utility function (a mathematical expression that assigns a value to all possible choices facing an economic agent or consumer) and where utility is defined as a measure of the relative satisfaction from or desirability of a consumption of goods. On an individual basis this idea works. Most of us will choose an outcome that maximises our own utility function when presented with a set of choices but the problem is that not everyone's utility function is going to be the same, and so when we are dealing with a group dynamic, what I would think the rational choice would be will probably be slightly different from others in the group. Nevertheless, this fact doesn't seem to stop people talking about *their own* rational choice as being the preferred outcome.

I have been in meetings so many times where a group of us are discussing the markets when one person says something that, to him, is a rational outcome and analysis of the situation but yet fails to address how the market as whole may react. I constantly ask myself why people do this. Why would you want to analyse something like a market that is such an obviously dynamic crowd from the point of view of one's own perspective? What I mean is, why would someone think that his own personal opinion of the future is going to be shared by the majority of the market's power (which obviously has to happen if the market price is going to move in that direction) when he should know that the market is made up of many diverse people with many diverse opinions over many different time scales. If I think that the economy is going to slow down and that it is going to be good for bonds, then what makes me think that (a) the majority of the market's power will think and act in the same way or that (b) a slowing economy will actually turn out to be good for bonds. The market's power here refers to the net demand or supply in the market. As was discussed earlier, it is not the amount of times that you are right but how much you are right when you are right and how much you are wrong when you are wrong that is the important factor in long-term trading or investing success. So in this context, if the majority of the market in *number* agreed with your own opinion but that the *buying power* in bonds of that majority turned out to be smaller than the *selling power* of the minority, then the price of bonds would not, in fact, increase. Could it be that the analyst in question thinks that someone else will hold the same opinion as him or her, and will be able to move the market "their" way because of the market power they have. Alas no. The motivation for people thinking about the markets from their own personal "rational" beliefs comes from a conscious or sub-conscious thought process that says "look, this is what I would think and do and so therefore this is the sensible course of action for the market to take". The point is that the belief in the market to react in the way that the individual rationalises the market is paramount in the thought process. There is a forcing of *personal will* on to the market and when there is a forcing of personal will on to the market then the person involved becomes a stakeholder in that *will* becoming reality. This personal will is then the "view" to which the person becomes wedded, and the cognitive dissonance process takes over.

SUMMARY

This chapter has revealed what are, in my opinion, the real, underlying, primeval forces behind the movement of prices in financial and commodity markets. Human beings drive the markets. If all the human beings left the stock exchanges and trading desks, do you think the markets would trade by themselves? No. Prices move because human beings are interacting with each other to trade and when prices of the goods human beings trade start moving, human emotions become the ultimate driving force behind decision making. It is human emotions, therefore, that are the engine of *volatility* in the markets and we can now move on to examine the nature of that volatility in more detail.

3

Volatility Defined

“One man’s freedom fighter is another man’s terrorist.”

ONE MAN’S FREEDOM FIGHTER

Volatility is arguably the least understood aspect of the financial and commodity markets. Mention the word volatility to some people in the markets and they would run a mile, as for them volatility is a bad thing. Mention it to other market participants and they will love it, because for them volatility is a good thing. We all talk about volatility as if it has a universal meaning but in actual fact it means many different things to different people and, depending on the way it is measured, it can have many different uses.

I can never understand people who work on the fund management side of the investment business when they say, and I have heard this a great deal over the years from some otherwise intelligent people, that they don’t like volatility and that volatility is a bad thing. The financial press compounds this perception that volatility is a big, bad, nasty beast when we see headlines like “Market Braced For More Volatility” or when the personal finance column advises people to “sit tight and ride out the volatility”. If they actually sat back and thought for a minute about this rubbish they would surely realise that it is this very volatility that they so deride that is responsible for the movements in the markets, and therefore responsible for making returns for their funds! What they are really saying of course is that they don’t like *downside* volatility in the stock market. That is, with most of these big mutual funds being long only in nature and the financial press assuming that the majority of retail investors are long only (which is probably a fair assumption by the way) they don’t like it when stocks go down. Being long only means that either they do not want to or, in the case of fund managers, do not have the mandate to “short” the markets and try to make money from falls in stock or asset prices.

Perhaps more than anything else this mass mandate of long-only pension, insurance and mutual funds has created a bullish bias in our culture where we are so skewed towards stock markets going up being “good” and stock markets going down being “bad”. Stock markets going down are not a bad thing at all for the long-run health of an economy if the economy has become overly imbalanced. Why is it that the popular perception of talking up a company’s stock price (having bullish analysis) is allowed but talking down a company’s stock price (having bearish analysis) is, in some way, very wrong? Hedge funds and other absolute return funds engage in short selling in order to seek returns for their clients’ money but, for some in the investment management industry, people that sell short are considered to be as detestable as rapists or murderers. One of the initial arguments a previous boss of mine in a long-only fund manager put forward for not wanting to engage in stock lending (where the fund lends out its stocks for a small return when they would otherwise just be sitting there earning nothing) was that he did not want to lend the stock to people who engage in short selling because it is morally wrong. The funds could have been making many millions of pounds extra in return for an extremely low risk, but because the boss thought that short sellers were in some way being “disloyal” to the industry, he didn’t want to take advantage of almost free money. You really couldn’t make this stuff up! Not only is that decision suboptimal but I would venture that it could be against a fund’s fiduciary duty to maximise the returns on client money.

Until hedge funds started gaining in popularity and became accessible for many more investors in the 1990s, the world and the financial press were conditioned to think of things going up as good and down as bad. The very word “crash” conjures up negative connotations but for some people who are short the market in those situations a crash is a very, very good thing. A downturn in stocks is usually associated with an upturn in measured volatility, but it isn’t always the case and, indeed, a rise in stocks is very often associated with an upturn in measured volatility. We didn’t hear too many fund managers complaining about volatility in 1995 through to 1997 even though the VIX index (the Chicago Board Options Exchange SPX Volatility Index, which reflects a market estimate of future volatility, based on the weighted average of the implied volatilities for a wide range of option strike prices) almost trebled. Oh no! That would be because the Dow Jones Industrials Average more than doubled in that period!

Volatility equals opportunity. Volatility moves up when prices are moving regardless of direction and when prices move in any direction that surely equals opportunity to profit from that move. If the stock market goes up on rising volatility you can buy more or you can make sure you are not short. If the stock market goes down on rising volatility you can sell short or, if you can’t do that, sell some of your long holdings. Either way it is an opportunity for you to make or save money. And that is what fund managers are paid to do for goodness sake! I’m sure a lot of long-only fund managers (definitely a lot of fund managers I have met in my career) would love a slow steady upward trend in stock prices where they can just watch the dividends roll in with a

little steady appreciation in capital value. Unfortunately the world does not work like that and we all have to accept volatility cycles as a fact of life in the markets. What makes the difference, though, is how we think of volatility – whether it is something to be avoided or something to be embraced. Forecasting markets is impossible but anticipating movements and managing risk is very possible. Therefore our attitude towards volatility is vital in determining the success of our overall game plan.

EQUILIBRIUM PRICE

The economics profession and certainly the academic economists love to talk about equilibrium, and therefore lesson number one in economics revolves around supply and demand curves and the equilibrium price. We are taught that when supply and demand are in balance in a market then the price is at an equilibrium level. Furthermore, when there is a change in demand, a change in supply, or both, then the market will find a new equilibrium price level at which to stabilise. This is the grandfather of all laws in economics, the most important law from which so much more, if not all, of economics stems.

Figures 3.1 and 3.2 show the classical stylised supply and demand curves (lines really, but economists still insist on curves!) and what would happen when there is a change in supply and a change in demand.

Figure 3.1 shows the classical supply and demand chart that we are all taught in our first lesson on economics. The vertical axis shows the price of a good or service and the horizontal axis shows the quantity of a good or service. The demand curve D slopes down from left to right because, as the quantity of a good decreases (becomes scarce), the price of that good will go up as more is demanded from buyers




Figure 3.1 Supply and demand




Figure 3.2 Supply and demand change

(all other things being equal, or “*ceteris paribus*” as they say in the economics ivory tower). The supply curve *S* slopes up from left to right because, as the price of a good increases, more is supplied by suppliers. The point at which the supply and demand curves intersect is the equilibrium price level – the price level where supply and demand are in balance and the amount supplied is the same as the amount demanded.

So far, so good. However, what if there is a change in supply or demand? The *reason* for the change in supply or demand is not important only that there is a shift in the quantity supplied and (or) demanded.

Figure 3.2 shows what happens when there is a change in the supply and demand dynamic. Let’s suppose first of all that there is an increase in demand. It is vitally important to realise that this increase in demand could have occurred for a multitude of reasons but that these reasons are not relevant to the analysis. The analysis here is that demand has increased. That is all we need to know. Due to this increase in demand the demand curve shifts to the right from *D* to *D1*. More is being demanded and, if there was no change in the supply curve, then the price of the good would move from *E* to *E1* where the new demand curve *D1* intersects with the existing supply curve *S*. More is being demanded, so more is being supplied (on the original curve at *E1*) and the price has increased to *E1*.

Now let us suppose that, subsequent to this increase in demand, there is a separate event that is a reduction in supply. Again, the reasons for this reduction in supply are not relevant, only that it has occurred. The reduction in supply causes the supply curve *S* to shift to the left from *S* to *S1*. Less is now being supplied (*E1* to *E2*) and, if there were no change in the demand curve, then the price of the good would increase from *E1* to *E2*.

So here we have had an initial increase in demand that caused the price to rise and then a reduction in supply that also caused the price to rise.

In this way the supply and demand dynamics interplay in a market for a good or a service and, to stress, in this analysis we do not need to know the *reasons* for the supply and demand changes. The reasons for the change in supply or demand would add nothing to this specific analysis. So if this is how supply and demand interacts in the markets for goods and services, then is there anything different to how the supply and demand dynamics interact in the financial markets? The answer turns out to be both yes and no.

From a theoretical perspective one could argue that there is, in fact, a very big difference in the actual shape of the demand curve when it comes to financial markets – rather than slope down and to the right, it actually slopes up and to the right. Rather than a lower price leading to higher demand it has been argued that, in the financial markets, it is actually the case that people demand more as the price goes up. This makes sense when we consider that it is a function of crowd behaviour and the herding instinct that predominate in the financial markets. So-called bubbles do happen in the financial and commodity markets, of that there is little doubt, and one of the reasons they happen is because of this upward-sloping demand curve phenomenon. The funny thing about this, though, is that our feelings towards a share price that is going up is something along the lines of “its going up so I want some but I want to wait for a pull back so I am getting it at a cheaper price than it currently stands”. When the share price doesn’t come back and continues higher we feel even stronger feelings of wanting to buy some shares and it is this psychology that helps the trend of the market along. For more on the subject of upward-sloping demand curves I would recommend the reader to look up www.socionomics.net. More on socionomics later.

For the purposes of this discussion, and to emphasise the point I want to make regarding volatility, the most basic supply and demand analysis, in my opinion, offers absolutely no difference between the philosophy, mechanics or outcomes. That is, the underlying dynamic of a market in any product, be it a real world good or service or a product in the financial markets, is the supply *of* and the demand *for* that thing and it is when supply and demand match up that the price of that thing is in equilibrium. By the same token, of course, when the supply and demand for that thing do not match up, then the market is in a state of disequilibrium and that is when prices move.

So when we look at a time series chart of the price of the US Long Bond, the Dow Jones Industrials, the share price of Google or the British Pound exchange rate, and we notice that there are times when the price moves in a range bound fashion as well as times when the market trades in a strong trending fashion, then what we are seeing is a market moving from a state of equilibrium to a state of disequilibrium, and so on and so on. When a market is in an equilibrium state the price is relatively static and supply and demand are matched. Although in this state the market is balanced in terms of supply and demand, it is in this state that we should think of the market as a whole having quite an uncertain psychological make up because neither the buyers nor the sellers are dominating. When a market is in a state of disequilibrium then the price is moving away from this equilibrium level and it is this state that we can

think of the market having a more certain psychological make up because either the buyers or the sellers are dominating. It does make me laugh sometimes when market commentators say something like “the market is falling because it hates uncertainty”. The market is *always* in a state of uncertainty. It is always in flux. What is true, though, is that the degree of that uncertainty changes dynamically due to the shifting dynamics of supply and demand, and this is what trading regime analysis aims to capture.

It is in this way that the basic laws of supply and demand underpin the movements in financial markets and analysing this relationship is, in fact, what technical analysts do all the time – identify what is happening to the relative supply and demand mix in the market. The reason for the change in demand or the change in supply is not important for the technical analyst, only that it is occurring. So what, in effect, the technical school is analysing is the most basic of fundamental principles (supply and demand) applied to the actual price or share price of the market or share in question. In a nutshell, technical analysts are applying the basic fundamental economic analysis directly to the tradable price of the market. Rather than conduct an economic analysis of supply and demand on something else in order to arrive at a conclusion about the market price, the technical analyst goes directly to the market to find out what message it is giving. The totally ironic aspect to the modern thinking of market analysis is that hardly any economists or fundamental analysts acknowledge the fact that the laws of supply and demand in the financial markets are as important as they are in the so-called “real economy”.

If we can agree, therefore, that it is the basic laws of supply and demand that drive the price movement in the financial and commodity markets, then what we are saying is that it is the basic laws of supply and demand that drive the price **volatility** in the financial and commodity markets. If the volatility of financial and commodity markets are driven by the dynamic interplay between supply and demand then, if we can identify when that dynamic interplay between supply and demand is changing, we can surely identify and anticipate changes in volatility regardless of the direction of the price. And anticipating changes in the volatility of the markets will enable us to anticipate if the market is more likely to be trading in a range or trading in a trend. Finally, if we can do that then we can pick the appropriate weapons for the trading or investing battles ahead.

VOLATILITY DEFINED

There are various measures of volatility, and depending on how it is measured it can reflect different aspects of price action. Whether a market is in a range-trading regime or a trending regime does depend on volatility but the nature of the price action can mean that, depending on which measure you use, the volatility might behave differently from what you expect.

We might start out by asking ourselves the question of what volatility is.




Figure 3.4 Volatility

The image we have in our heads of something being volatile is usually something that goes up and down sharply in a mean reverting fashion. That is, it goes up and down in a defined, almost horizontal range. However, what we find when we examine the actual calculations of volatility a little more closely is that the picture we should be thinking about is probably something closer to Figure 3.4; that is, something that moves sharply, in a *trend* like nature, away from previous values.

CALCULATION OF HISTORICAL VOLATILITY

Historical volatility in the financial market context is a calculated measure of the change in price of a financial instrument over a given period of time. It is calculated as the standard deviation of logarithmic price changes expressed as an annualised percentage. Therefore, 10-week historical price volatility represents the annualised standard deviation of relative price changes calculated by analysing a sample of the 10 most recent weekly closes. It can therefore be thought of as the standard deviation of the continuously compounded return over the time period in question. Note that historical volatility is not a forecast of future volatility; it only quantifies the level of volatility that has been observed.

Let's look at an example using five daily closes of an exchange rate.

Figure 3.5 shows the New York closing exchange rates for a currency pair over one week. To calculate the historical volatility we take the current day's closing price divided by the previous day's ($X/(X - 1)$) and then we take the logarithm of that number. We then take an average of these logarithms in order to find the deviation that each logarithm has from the average. We then square each deviation before totalling them. The variance can then be calculated, which is this total divided by

The *Oxford English Dictionary* states that volatility is “the quality, state, or condition of being volatile; readiness to vaporize or evaporate; tendency to be readily diffused or dissipated in the atmosphere, especially at ordinary temperatures; tendency to lightness, levity, or flightiness; lack of steadiness or seriousness; and adaptability for flight”.

Dictionary.com defines volatility as “tending or threatening to break into open violence; explosive” and “(of prices, values, etc.) tending to fluctuate sharply and regularly: *volatile market conditions*”.

Investopedia defines volatility as “a statistical measure of the tendency of a market or security to rise or fall sharply within a period of time”.

By studying these definitions we get a sense that volatility is used to describe something that has a lack of steadiness in that it fluctuates sharply and regularly within a period of time. Using these definitions, a time series chart for a market described as having a lot of volatility might look like the diagram in Figure 3.3.




Figure 3.3 Volatility

This diagram seems to depict the descriptions of volatility that come from the dictionary sources. The market moves up and down in a sporadic random fashion within a time period and this, I would guess, is the picture most people imagine when they think about something being described as *volatile*. But what about Figure 3.4? This diagram also fulfills the description of volatility as described by the definitions because this market “rises or falls sharply within a period of time”.

It is important to note that the definition in *Investopedia* alludes to a market that rises *or* falls sharply, not rises *and* falls sharply that most people would tend to think about when we talk about volatility. This may seem like semantics but, to me, this is a key determinant in our understanding of what volatility is and how to identify it. Volatility is not only referring to something that fluctuates sharply up and down but is also referring to something that moves sharply in a sustained direction.

	Closing Price	$x / (x-1)$	Logarithm $x / (x-1)$	Deviation	Deviation Squared
Monday	1.55				
Tuesday	1.5465	0.99774194	-0.002261	0.002422	0.0000059
Wednesday	1.5533	1.00439703	0.004387	0.004226	0.0000179
Thursday	1.5587	1.00347647	0.003470	0.003309	0.0000110
Friday	1.551	0.99505999	-0.004952	0.005113	0.0000261
			0.000161		0.0000608
				Variance =	0.00002027
			Standard Deviation =		0.004502753
				Days per year	252
				Square root of days	15.8745
			Historical Volatility =		7.15%

Figure 3.5 Historical Volatility

the number of data points minus one. In this case we have five days of the week but, because we are calculating the standard deviation of the *changes* in the prices, we have only four data points. Therefore, in order to get the variance we divide the sum of the deviations squared by 3.

Once we have this number we can find its square root, which is the standard deviation. It is critical to note that this is the standard deviation of the *changes* in the prices and not the standard deviation of the price series itself.

Finally, once the standard deviation of the changes in the prices has been found we can multiply that number by the square root of the number of days per year to find the annualised historical volatility of the time series. In this case it turns out to be 7.15%.

So the calculation of historical volatility has at its core the concept of deviation from the mean or deviation from the average. But it is the standard deviation of the *changes* in the prices that it is calculating. Therefore a market trading in a range could, if it was experiencing sharp up and down movements, have a high measure of historical volatility like that shown in Figure 3.3. On the other hand, if the market was experiencing a strong trend, as depicted in Figure 3.4, but the changes in the price series within that trend were quite uniform, then the measure of historical volatility could be quite low.

Therefore our perceptions of what constitutes a volatile market and what doesn't, can differ a lot from the actual calculations of historical volatility.

Figure 3.6 shows the EURUSD exchange rate using the weekly time fractal (the bar chart of *weekly* high, low and close data) and the 20-period historical volatility that has been calculated from that data (see Appendix 1 for an explanation of time fractals). We can see that the market was in a downtrend until late 2000 when, at point A, it started to trade in an up and down fashion until point B when it subsequently broke out into a trend. If we just looked at the price action between point A and point B then we observe that the market goes up about 11%, down about 11%, up about 9% and down about 8%, all within the space of just over a year. Does this constitute a volatile market? In our heads we would think that it does because we are conditioned to think of volatility as something that means a market that goes up and down a lot. However, if we look at what actually happens to the measurement of historical volatility then we can see that it actually goes down a great deal during this period.



Figure 3.6 EURUSD weekly – historical volatility

Given the fact that historical volatility is, as the name suggests, very much a lagging indicator of volatility, then the blip up in volatility after point A is mostly referring to the period just prior to point A when the market's downtrend was in full force. Once the price range is established, though, historical volatility starts to drop, and drop a lot. So the idea we have of a market going up and down a lot as being volatile is not strictly correct. It is our perceptions of what volatility means that makes us think of something oscillating about a lot and, as we shall see later, that something oscillating about a lot will probably occur when volatility has already risen and is already high. Correspondingly, it is when volatility is low that a market will break out into a trend and the price action will not oscillate. This we can see just after point B when the market breaks out into a trend and volatility starts to rise up to point C. Then after another period of oscillating price action, when volatility declines, the market's up trend continues with volatility rising as well.

Looking at Figure 3.6 we could conclude that there appears to be a regular cycle to volatility in that it very broadly reverts back to some sort of average level after a period of divergence. It goes up and down in what appears to be a regular fashion. Does that mean it can be "predicted"? I would argue probably not, but, as we shall see, not being able to "predict" something does not mean that a cycle turn in volatility cannot be anticipated.

THE ONE MARKET CONSTANT?

For liquid markets, the one constant truth is that they go from periods of high volatility or range expansion to periods of low volatility or range contraction over and over again. In other words they fluctuate between periods of equilibrium prices when levels of supply and demand are similar, and periods of disequilibrium prices when levels of supply and demand are skewed one way or the other. Whether you are a fundamental or technical analyst, whether you believe in the Elliott Wave Principle or the capital market asset pricing model, or whether you believe in efficient or non-efficient markets, the one underlying constant in liquid market behaviour is that markets go from periods of range trading to periods of trending.

From Figure 3.6 we can see that volatility is low at the start of trends and high at the end of trends, so if we can accept that volatility generally peaks *after* a trend in the price and generally troughs *before* a trend begins, then surely we must think of volatility, when it starts to move higher, as being a friend rather than a foe. Yet still the popular perception is of volatility being a bad thing.

What people tend to forget is that volatility is mean reverting. It goes up and down within a generally well-defined range and anything that does that consistently should provide many high probability opportunities. Something that is generally mean reverting will move in cycles that can give the analyst a very good idea of where relative value is. The issue, of course, as with any mean-reverting indicator or process, is the use of the word "consistently". The reason that my trading or investing

bias is towards trend following is because I have learnt over the years that mean-reverting indicators such as relative strength indicator, stochastics and other trend exhaustion techniques, can go through periods of remaining overbought or oversold for a long, long time. As my old boss said to me once when he was paraphrasing John Maynard Keynes, “Murray,” he said “the market can remain overbought or oversold longer than you can remain solvent.” In other words, don’t rely on something mean reverting just because of what it has done in the past. This in fact was the seed of my motivation for trading regime analysis because it is obvious that mean reversion techniques will work very well in range trading markets and work terribly in trending markets. Equally, trend following techniques will work brilliantly in trending markets and work terribly in ranging markets. If we can identify the likely trading or volatility regime we are going to be in, then we can pick whether we want to trend follow the market or use a contrarian, mean-reverting style.

So what of volatility then? Is that not just another mean-reverting indicator that can go through long periods of being low and long periods of being high? The honest answer is “yes” it can but, as we shall see, that does not alter the fact that we can add value by adjusting our trading style mix or position size to ready ourselves for the inevitable change in trading regime conditions.

Volatility mean reverts, but it can do so in a very fast and furious way and the quantitative analysts call this *clustering*. It was the mathematician Benoit Mandelbrot who noted that in volatility there is a tendency for big moves to be followed by other such big moves and, similarly, for small moves to be followed by other such small moves. That is, volatility tends to display a correlation with previous values and this has spawned a lot of research in the field of forecasting volatility using models like ARCH (Auto Regressive Conditional Heteroskedasticity) and GARCH (Generalised Auto Regressive Conditional Heteroskedasticity). While these models try to fine tune the forecasting of volatility levels, their core idea is that volatility is dependent upon past values of volatility and that volatility tends to revert to an average (or mean) rather than remaining constant. A whole book could be written on volatility models using ARCH and GARCH so I shall leave it to the reader to explore these further.

Figure 3.7 shows the VIX Index (showing the implied volatility in the US stock market) in the weekly time fractal and we can see that the index can go through periods of grinding lower in a slow, stable fashion only to blast higher from time to time in a fast, dynamic move. In layman’s terms, this is what is meant by volatility clustering and it follows that “forecasting” volatility can be extremely perilous.

The main point to note is that volatility clusters lead to a mean reversion process that can be very fast and furious and, at first glance, the analyst will think that it would be impossible to take advantage of that reversion back to the mean. However, this is when an appreciation of fractal trading comes into its own as it is still possible to take advantage of a quick move up in volatility if we trade one fractal (or time period) down from the volatility analysis. Using option contracts can also allow the investor or trader to benefit from this fast explosion in volatility.




Figure 3.7 VIX weekly – volatility clustering

IMPLIED VOLATILITY

Historical volatility calculates the volatility of a market. Implied volatility refers to the volatility level that is implied for the market via the pricing of options. Implied volatility and historical volatility will tend to be somewhat similar but can differ from period to period, and it is these differences that relative value option traders look for when they are searching for opportunities within the options market. For the trading regime analyst, however, the difference between implied and historical volatility is not that important, only the absolute levels of both. The implied volatility can be thought of as the market's guess or expectation of future volatility and, as such, the logical thought would be that it might affect how people behave in the underlying market. The underlying market refers to the market on which the options are traded, so if we were trading options on the share price of Microsoft then the underlying market would be the share price of Microsoft and if we were trading options on the S&P 500 futures market then the underlying market would be the S&P 500 futures market.

In actual fact it is probably the underlying market's behaviour that affects the pricing of implied volatility more than the other way round, because despite a traded market existing in implied volatility the reality shows that it only rarely goes very far away from historical volatility. Options market makers will buy and sell volatility

volatility, but also that it by no means captures all changes in trading regime using the logical methodology of declining volatility giving rise to a range-trading regime and advancing volatility giving rise to a trending regime.

RELATIVE VALUE VOLATILITY ANALYSIS

There are essentially two ways of evaluating volatility: absolute and relative value. Absolute volatility analysis is what this book is concerned with because it examines the probability of absolute volatility occurring in the underlying market. That is, we are trying to anticipate whether the underlying market price will be more or less volatile in the next period. The other way of evaluating volatility is the relative value approach where the analyst is more concerned with the relationship between historic (or realised) volatility and the level of volatility being implied by option pricing. This approach is undoubtedly essential in analysing the relative attractiveness of options but for an analysis of anticipating absolute volatility it is patchy at best.

It is extremely attractive to think that implied volatility should act as a form of anticipator of future volatility or trading regime itself. So if implied volatility was higher than historic volatility, then it would be signalling that a trending regime was in force, but if implied dropped below historic then it is signalling that a range-trading regime was developing. This is, after all, what the implied volatility should be discounting in its pricing. Given the way the emotional human market behaves, it might be logical to assume that, as a market enters into a strong trending phase, the level of implied volatility will be above the level of realised volatility and rising. That is, as the market trend develops, the level of fear that the trend will extend grows and implied volatility becomes much higher than historic volatility. On the other side of the coin we might logically expect that, as a market goes through a range-trading phase, the level of implied volatility will be lower than historic volatility because the market becomes complacent that the range will continue. However, while it is extremely attractive to think that this dynamic works, in practice it would be wrong to make that assumption.

There are undoubtedly occasions when the implied to historical volatility spread correctly anticipates a trading regime, but there are equally many times when it does not.

Figure 3.9 shows the GBPUSD exchange rate in the weekly time fractal and the difference between 3-month realised and implied volatility. When the area chart is below zero, implied volatility is below realised and can be considered to be “cheap”, whereas if it is above zero it can be thought of as being “expensive”. We can see that there is some evidence to suggest a link, in that when the market is ranging the implied volatility can become quite cheap and when the market is trending it can become expensive. However, the link is patchy and so, although relative value volatility analysis is undoubtedly essential for relative value option traders, its use to

(implied volatility) from and to each other all day long but they will be guided in pricing the level of volatility by the options-pricing models that they use, the most popular of which are based on the Black–Scholes Model.

Figure 3.8 shows the EURUSD exchange rate on the weekly time fractal in the upper portion of the chart and the 1-month implied volatility in the lower portion of the chart. This is the market level of implied volatility for an option with a 1-month tenor. We can see that at point A the underlying EURUSD exchange rate ended its preceding downtrend, and that corresponded to a top in the implied volatility of the exchange rate. Then as the EURUSD rate traded in a distinct range until point B, the implied volatility generally declined before turning higher into the move up to point C. Notice, however, that as the EURUSD exchange rate continued its strong trending action after the range immediately post point C, the implied volatility went mostly sideways. So what we can gather from Figure 3.8 is not only that implied volatility tends to give a more coincident indicator of trading regime than historical



Figure 3.8 EURUSD weekly – implied volatility




Figure 3.9 GBPUSD weekly and Relative value volatility

gauge the potential trading regime of a market should be used only in conjunction with other analysis techniques.

STANDARD DEVIATION

When we are measuring the historical volatility of a market we are measuring the standard deviation of the *changes* in that market's price over a certain time period. It is absolutely critical to realise that this is not the same as measuring the historical standard deviation of the market prices themselves, which is a measure of the *dispersion* or *spread* of the market (not the changes in market prices) over a certain time period.

The standard deviation is the square root of the variance in a time series, with the variance being defined as the average of the square of the deviations from the mean. As a simple example, consider a stock price that, over the course of one week, had a closing price series of 10, 12, 15, 11 and 8. Figure 3.10 shows the steps involved to arrive at the standard deviation. The average of the closing prices is 11.2. The sum of the square of the deviations comes to 26.8. This number divided by 5 (the number of data points in the time series) gives the variance of 5.36. Finally, the square root of the variance gives the standard deviation of 2.315.

Figure 3.11 shows the statistics for two time series of stock closing prices. The stock at the top has experienced a week where the market has fluctuated in a narrow

range between 8 and 12, and due to this fact the standard deviation of the time series is low, at 1.789. However, because the *changes* in the stock prices have been quite high *and* in different directions, the historical volatility is a massive 657.08%.

The stock at the bottom, however, has experienced a week where there has been a strong uptrend in price with the market moving from 10 to 24. Due to the fact that the spread of prices is greater, the standard deviation is much higher than the other stock at 8.198. However, notice that the historical volatility has dropped dramatically to 78.97%. This has happened because, even though the numerical changes in the stocks are the same (2, 4, 4, 4), the changes here are in the same direction.

Therefore, when we come to think of trading regime analysis, and try to work out whether the market will be in a range-trading regime where the market prices will be confined to a range, or in a trending regime where the market prices will be spread out, it makes much more sense to think that an analysis of the standard deviation of the market prices themselves will be more useful than an analysis of the standard deviation of the *changes* in those prices. The speed of the changes in the price of a market is not as important as the distribution of those prices when we want to analyse the current trading regime and the potential future trading regime.

VOLATILITY VERSUS PRICE EXPANSION AND CONTRACTION

This is perhaps the most important section of the book because it explains that my definition of volatility in relation to trading regime analysis is referring to the expansion and contraction of price *ranges* rather than the speed of change in the prices. This is what confuses people when it comes to trading methods such as volatility breakouts and such, because that little word volatility can actually be referring to different calculations. It is the expansion and contraction of price values themselves, not how quickly the changes in those price values within that expansion or contraction are taking place – that is, the manifestation of trends or ranges in market price action – and, moreover, it is the human behavioural nature of markets that are the underlying driver of this volatility cycle. When we calculate volatility using standard deviation rather than historical volatility, then we can see that there is a much more rhythmical cycle to it. It highlights periods when the price is obviously contracting in a trading range but, whereas the standard deviation is declining, the historical volatility measure is not declining and could in fact be rising.

Figure 3.12 shows the bar chart of the euro exchange rate with the US dollar using the weekly time fractal and, in the bottom chart, the associated historical volatility and standard deviation statistics. The thick line in the bottom chart is the standard deviation of the last 20 periods (using the scale on the left) and the thin line is the historical volatility of the last 20 periods (using the scale on the right). We can see immediately that there is a difference in the calculations and that there are times when both lines are going in opposite directions. If we zone in on the shaded area, which is the period between the start of 2002 and the middle of 2003, we can see

		Closing Price	Deviation	Deviation Squared	
	Monday	10	1.2	1.44	
	Tuesday	12	0.8	0.64	
	Wednesday	15	3.8	14.44	
	Thursday	11	0.2	0.04	
	Friday	8	3.2	10.24	
	Average =	11.2		26.8	
			Variance =	5.36	
			Standard Deviation =	2.315	

Figure 3.10 Standard Deviation




Figure 3.11 Trading Regime and Volatility




Figure 3.12 EURUSD weekly – historical volatility versus standard deviation

that, as the euro started to trend higher against the US dollar, the standard deviation measure picked up, and did so quicker than the historical volatility measure. When the uptrend had come to an end and the market entered a period of sideways movement, the standard deviation measure then started to drop very quickly, whereas the historical volatility measure stayed the same and in fact started to increase a little. This is a classic example of how using a measure of historical volatility as classically calculated is not necessarily the most relevant indicator for trading regime analysis. The market is obviously trading in a range-trading regime here and yet the historical volatility measurement is telling the analyst that the market is still quite volatile. The standard deviation measurement, on the other hand, is telling the analyst that the market is trading in a range. So we have a situation where the market is obviously trading in a range and yet the dispersion of the changes in price is still quite high. When we think about it, this is actually a logical situation because a range-trading

is on the right-hand scale while the standard deviation scale is on the left. The vertical lines on the chart indicate where the standard deviation has reached what could be thought of as historically low levels and we can see that the drop in the standard deviation to these low levels usually, not always but usually, precedes a strong trend like move in the price of the US Dollar Index. This is simply what we would expect given the calculation properties of the standard deviation. The standard deviation is calculating the distribution of price values over a specified time period and so if the standard deviation is low then it is telling the analyst that the distribution of price values has been occurring in a tight range. That is, the market has been range trading. Given the mean reverting nature of the standard deviation then, after a period of range trading, it is reasonable to expect the market to enter a period of higher standard deviation and that usually means a trending period. Sometimes, of course, the resulting price action is not a strong trend and the analyst has to be open to this possibility. However, the *probability* lies with the fact that after a period of price range contraction comes a period of price range expansion.

In addition to a low standard deviation preceding a trending move in the market, a high standard deviation can precede a range-trading move in the market. On the left-hand side of Figure 3.13 we can see that the US Dollar Index was in a strong downtrend having come from a low standard deviation, but then around July 2002 the standard deviation had reached very high levels and began to turn over. There then followed a number of weeks of range trading with the standard deviation subsequently collapsing. Once the standard deviation had reached low levels again it was a precursor to yet another strong trending move down. So, low levels of standard deviation are usually followed by a trending market and high levels of standard deviation can be followed by a range-trading market. Sometimes a high standard deviation reading that has come about due to a strong trend in the underlying market can be followed by a period where the standard deviation is falling, but the market does not follow a strict range-trading pattern as defined by a very narrow range. However, when the standard deviation is falling it usually coincides with choppy price action.

In this regard, a high standard deviation reading does seem to be a good trend exhaustion indicator. Look at the times when there is a spike up in the standard deviation. When it starts to turn over it usually signals at least a temporary end (or exhaustion) of the previous trend in the market.

TREND EXHAUSTION

It is human nature to move in fits and starts. The Elliott Wave Principle has defined this as three moves up followed by two moves back. This is how progress is made and it is a very natural phenomenon. Therefore trends will go through periods of replenishment before moving off again in the same direction. These periods of replenishment can be thought of as periods of trend exhaustion points, at least in the context of that part of a bigger trend. Sometimes of course the trend will exhaust

regime can involve – not always but sometimes – a wild fluctuation in the market price that is still bound by an upper and lower limit.

After this period of range trading, which ended around November 2002, the euro again took off versus the US dollar, rising in a trend-like fashion from around 0.96 to around 1.08 over the next few months. The standard deviation measure that had collapsed down to historical low levels while the market was trading in a range, exploded higher again as the market entered the period of trending. However, look at the historical volatility statistic! It was actually *declining* as this strong trend was taking place and, again, this is another classic example of why it is the standard deviation of the price values, not the changes in those values, that is important in analysing volatility conditions in the context of a trending or range-trading market regime. Here the market is obviously in a strong trend but the crucial aspect is that the changes in the price during that trend are taking place in quite a uniform fashion, and this means that the measure of historical volatility will fall. Notice too that, once again, after the trend had ended and the market entered a period of range trading, the standard deviation fell but the historical volatility actually rose.

An appreciation of the difference between standard deviation and historical volatility is vitally important when we come to analyse the current and potential future trading regime of a market.

Figure 3.13 shows the US Dollar Index (thin line) in the weekly time fractal and the 20-week standard deviation (thick line) from 2002 to 2007. The US Dollar Index




Figure 3.13 US Dollar Index weekly and standard deviation

and rather than continue up in the same direction, after a while a new trend will be established in the opposite direction. Nevertheless the starting point of trend exhaustion is what we as trading regime analysts are interested in because it will help us to determine when, as the name suggests, a trend is approaching an end. However, when it comes to trend exhaustion identification, the mistake a lot of people make is that they assume that when the current trend “exhausts” itself a new trend in the opposite direction, or at the very least a sizeable move in the opposite direction, will naturally occur. This belief is a natural manifestation of our instinctive human desire to try to pick the top or bottom of the market. It is human nature to want to buy something because it has gone down a lot and it is human nature to want to sell something because it has gone up a lot. For this reason technical indicators that show overbought or oversold levels and methods that try to pick trend “reversals” are very popular. It goes against human nature to buy something that has already gone up or sell something that has gone down because our brain is focusing on the previous price action and is registering that a move has already taken place. Interestingly, this only happens at the beginning of a trend. If the trend turns out to persist, by the end of the trend the people who refused to buy it after it had made a modest up move because it “has already moved higher” are the same people who are scrambling to buy it later and, of course, this is the point when the last buyers come into the market and the trend exhausts itself.

Figure 3.14 shows the Dow Jones Industrials Index (thin line) in the monthly time fractal and the associated 20-month standard deviation (thick line) from 1960 until 1985. The index scale is on the right and the standard deviation scale is on the left. The period being analysed was categorised by a lot of sideways price action until the early 1980s when the Dow took off at the start of the greatest bull market run in history. We can see again that when the standard deviation is low then it usually, not always but usually, precedes a strong trend-like move in the market. Also when the standard deviation spikes up and starts to turn over it usually signals trend exhaustion and possibly some choppy price action. If we focus our attention on the period from 1977 we can see that from around the end of that year the market traded in a very choppy, range-trading fashion and this was associated with a persistent fall in the standard deviation. The standard deviation finally bottomed around March 1980 at the lowest level it had been at since the end of 1968. The Dow then rallied some 40% over the next 18 months and, following the subsequent setback when the standard deviation made a higher low towards the end of 1981, it marked the beginning of the massive bull run that lasted until 2000.

Now, at this point, the obvious question that the reader will be asking is what constitutes a low reading of standard deviation and what constitutes a high reading of standard deviation. I would dearly love there to be a cast iron rule that could consistently identify a high level and a low level of standard deviation but, while I address this question in Chapter 21 with the Trading Regime Indicator, the cold hard truth is that the standard deviation is essentially an unbounded mean-reverting oscillator and, as such, there is, in theory, no absolute level where it can go no further.




Figure 3.14 Dow Jones Industrials monthly and standard deviation

Well it can obviously not go below zero but you know what I mean. We can model and optimise as much as we like but, as with any model of past history, there will always be times when new historic highs or lows will be registered and cause the model to get out of kilter.

VOL OF VOL

The standard deviation itself proves to be a useful measure of when prices are constricted within a range or are reaching an exhaustion point having had a period of range expansion. What if we were to look at the standard deviation of the standard deviation itself? Would that add any more value to our trading regime analysis? A number of people have looked at the volatility of volatility (or volga as it is commonly known) but they have looked at it from the traditional standpoint of historical volatility and implied volatility. Given the fact that when analysing the trading regime of the market my preferred measure of volatility is the standard deviation, I wanted to look at the standard deviation of the standard deviation.

Figure 3.15 shows the standard deviation of the standard deviation of the EURUSD exchange rate using the weekly time fractal. We can see that, although there are times

with that, but the philosophy of comparative advantage worked and still works so well in practice because it really can be a case of everyone benefiting (at least *relative* to where they were before). The call centre worker in India today might well be made to work extremely long hours in extremely harsh conditions, but he is still better off relative to what he would be doing in the absence of the call centre. The work of the classical economists continually reminds me that economics is perhaps much more of an exercise in social philosophy than an exercise in hard, mathematical econometric science and, in this regard, the classicists are more aligned to the aspect of behavioural influences in economics which they regard as intertwined with the very fabric of society.

I was brought up in a traditional manner. Christened at birth, Sunday school and the odd trip to church over my childhood years were pretty standard for someone my age from where I come from. However, my parents were not the sort to indoctrinate a strict religious belief in me but guided me in my morals to give me a basic sense of what is right and what is wrong. So being left thankfully to develop my own free thinking and not to slavishly follow one particular dogma, over the years I, just like a huge number of human beings, have struggled with the age-old question of what role religion plays in life and whether there is a God or not. The more I have read, thought, experienced life and philosophised about this the more I have come to the conclusion that the overwhelming evidence, at least in my opinion, suggests that whether there is a God or not (whatever that means to whichever brand of religion or sect in question), life on planet earth evolved from bacteria in the seas and that the chain of evolution brought humankind from the apes. I believe that there are just so many similarities between humans and primates in particular that to deny a genetic link would seem to take denial to a new level. Now, whether some sort of creationist God was responsible for the process of evolution I will leave to the theologians and scientists to battle over, but suffice to say that I am a believer in Darwinian evolution and, from where I stand, we humans are essentially apes who wear clothes. It is my belief that we are, in a very real sense, a member of the animal kingdom.

A large part of this belief I may add has come from observing human behaviour in stressful situations and, if we human beings are evolved animals, then it makes total sense that we should exhibit tendencies in our social behaviour and psychological make up within those situations that are very similar to other members of the animal kingdom. It is well documented that we do indeed act in a very similar way to animals, especially when part of a crowd. In fact the very essence of crowd is core to our behaviour. We feel safer in crowds just as most animals do. The family unit is of course the basic “crowd” but we humans also feel good belonging to clubs, societies, organised religion, sporting teams and so on. African tribes are as good an indicator as any that human beings like to associate themselves with a group. When human beings are in groups the dynamics of psychology can change. If anyone has ever found themselves at a sporting occasion and being “carried away” with the emotion shown to their team by the massed supporters, then they know what I am talking about. People’s behaviour can change radically when in a group situation. Someone

(highlighted by the arrows) when the indicator is extremely low just before a trending regime and extremely high just before trend exhaustion, the concept of this “vol of vol” does not add much to the vanilla standard deviation.




Figure 3.15 EURUSD weekly and vol of vol

ANIMAL SPIRITS

The older I have become I have grown more and more intrigued with the big questions of life. Who are we and how did we get here? Why does history appear to, if not exactly repeat itself, as Mark Twain said, “rhyme”? I love these sorts of philosophical questions and, as a result, I have always been drawn to reading the classical economic philosophers. Men like David Hume and Adam Smith were great philosophical thinkers of their time and it makes me very proud as a fellow Scotsman to think that their writings and work gave so much influence to the modern world in which we now live. Their work, particularly Smith’s *Wealth of Nations* had a direct influence on the writing of the constitution of the United States of America in the 1780s and shaped the first great era of globalisation that was dominated by the British Empire. Globalisation is not, as many think, a new phenomenon! The East India Company was “outsourcing” in the 1600s for goodness sake, albeit with much more ruthless employment and trading practices than in the modern world. Come to think of it, perhaps some who work in the modern day call centres would disagree

could be as mild mannered and as polite as you like when away from the soccer ground, but when the match kicks off he turns into a foul-mouthed hooligan as the group of supporters takes on a life of its own. If everyone else is singing for the team then our man will find it almost impossible not to join in, and if the mob turns nasty and vehement towards the referee then again our man, no matter how hard he tries, will probably succumb to taunt the referee, just like everyone else. This is because we humans have a basic desire to belong to something and not to be excluded from the group. Not everyone is the same of course, and there are people in this world who are exceptions to this generalisation. However those people will still feel the same emotions as the rest of us when it comes to group situations; it's just that they have learned to control their responses to those emotions.

If we examine herding behaviour and look at the aspect of movement then we can get a sense of how groups and crowds move. Take a troop of chimpanzees, for example. They could be going about their daily business of grooming each other and lounging around when something scares or spooks one of them. Almost instantly there is a kerfuffle among the group and a flurry of activity before things eventually calm down. What about a herd of zebra on the African plain? The group could be grazing happily when they sense a lion or some other danger in the area, whereupon one of them starts to flee off with the others following. These animals could not go on fleeing or being in a state of distress forever as there is finite capacity for any animal in terms of stamina so, at some point, the frantic nature of the behaviour calms down and things return to a normal (or equilibrium) state. The human animal is no different. Take riots as an example. When groups of people riot the pattern is never one of constant and consistent battling with the police or looting. Instead the rioting and looting seems to go in cycles of intensity. The human animal needs to physically rest but also needs to mentally rest in order for the brain to absorb fast-moving events.

What if we had a measure of intensity during riots or a measure of the intensity of activity in herding animals? If we had such a measure then it might show periods of low intensity followed by periods of high intensity. Well, we do actually have such a measure for the human animal when their behaviour is part of a group, herd, tribe or crowd and it is known as market volatility! The volatility of a market measures the periods of activity and inactivity, of stress and relaxation that a group of animals, in this case human animals, go through in their behavioural cycles.

Images of the open outcry trading "pits" in the futures markets encapsulate this notion of the market being a crowd whose behaviour is influenced *en masse* by group behavioural dynamics. There are some legendary stories about the trading pits, some of which I am sure are urban myths, but the one that sticks in my mind is a shocking one which, if true, is a particularly raw evaluation of human crowd behaviour. The bond pit (where the US government bond future is traded) in the Chicago Board of Trade is massive. The size of a football pitch, it holds many hundreds of people who each day stand and trade futures with each other either executing orders for people off the floor in, for example, fund managers' offices or people trading on

their own account (known as locals). It is widely regarded to be the epicentre of the capitalist universe. Bond folk love macro-economic statistics as this is the lifeblood of how a lot of expectations on interest rates are formed and the grandfather of all the economic statistics in the United States is the Non Farm Payroll number. It wasn't always the case that the "NFP" number was considered to be the most important, as many older readers will remember when the money supply economic statistics were given that status, but over the 1990s and 2000s, as the macro-economic landscape changed from inflationary to dis-inflationary, the Non Farm Payroll number has been considered to be the bell-wether number for gauging the health of the US macro-economy. The Non Farm Payroll report measures the change in the number of paid US workers of any business excluding farm employees and such like. It also gives estimates of the average weekly earnings of employees, and these employment statistics are considered to be so important to the health of an economy because so many other economic variables feed off the basic fact of whether someone is employed or not. If you are employed then you can make long-term plans like buying a house, for instance, which then affects other sectors like firms that build fridge freezers. Or if you are employed you can spend some money on entertainment like going to football matches and going out to restaurants. All of these things, along with a billion others, affect the general make-up of the macro-economy but a great many of these factors are directly affected by the employment profile of the population. Therefore, suffice to say that bond traders regard the Non Farm Payroll report as the most important piece of macro-economic information they receive each month.

On this particular month there was the usual anticipation and excitement before the release of the number and the bond pit was, as always for NFP, jam-packed. On days like this it is very real drama with usually a perceptible pause in noise level a few seconds before the release of the number as all eyes in the pit are on the board showing what the number is. In fact these days the whole of the financial markets on trading floors and offices around the world seems to pause for a few seconds before the number comes out. A split second after the number is released there is a massive burst of activity and, depending on whether it was near to what people were expecting or not, this can go on for quite a while before exhaustion sets in. Now, there is a wide range of ages in the pits with some young, athletic and generally tall (in order to get an edge in the open outcry environment) people, and also some grizzled old timers who have been there, seen that and, as they say, got the t-shirt. Well one old timer on this day was in the throng of the crowd when the Non Farm Payroll report was released, and as this particular report was far away from what people had been expecting, the burst of activity and rush to trade was even more frantic than usual. Very unfortunately for our old timer, his heart decided to pack in just as the number came out and he collapsed there and then on the spot as his heart attack took hold. Now, in any other environment, if someone standing next to you collapsed you would most probably go to his or her aid, but somehow when we are in a crowd our behaviour can change dramatically. According to the story, this poor chap was left on the floor of the pit as traders, who had noticed him but were

too wrapped up in the emotions of the crowd and the moment to help, stepped over him as they continued to trade. According to one story teller, he claimed that traders whose trades over the next few minutes had gone the wrong way and accrued losses, stuffed their losing tickets into this poor man's jacket pocket because they thought, if he had died, then their losing trade would "die" with him. Personally I do not believe that part of the story, because futures pit traders are each identified by numbers and it would be very difficult to try to pass on your trades as belonging to someone else. Nevertheless, the story does shock because it reminds us that a mob in full fray has a collective mind and behaviour that no normal human being would comprehend outside of the crowd. I have it on good authority that this story is true but, even if it is not, there is a host of other anecdotes about the trading pits that involve physical fighting, verbal abuse and emotional breakdowns. Behaviour that, were it not for the crowd dynamic, would generally never be enacted by the individuals concerned.

There are services now where day traders who sit at home or in their office have the ability to listen to the noise that emanates from the open outcry trading pits. They do this because the noise level will fluctuate up and down depending on whether there is activity or not, with the assumption being that a noisy pit will be associated with a market that is moving and a quiet pit will be associated with a market that is not really moving very much. If anything encapsulates the fluctuating cycle of activity (trends) with inactivity (ranges), then this is it in real live stereo piped into your home office! At one time only the locals who were physically in the pit would have had this edge. There would be a flurry of excitement as some big orders came into the pit from, for instance, a large institutional player and the locals would jump on board to "coat tail" the order, meaning that if the order was to buy a lot of contracts then a local would buy one for himself in the expectation that the big order will move the price higher. Now the home trader can listen as the noise level rises with the order being recognised by the market participants in the pit, and it might give him or her an extra split second to place their trade. The mountain really *has* moved.

The intensity and loudness of noise and activity from an open outcry trading pit is correlated with how much the market is moving. So if we had a decibel meter machine to measure the decibel level from the pit and plotted it against the market price, it might look something like that shown in Figure 3.16.

The top chart in Figure 3.16 shows the US 30-year Treasury Bond future in the 5-minute time fractal over the course of two days' trading in April 2008. The bottom chart shows a proxy for a decibel meter chart and we can see that it moves up when the market is moving (in whatever direction) and moves down when the market is not moving very much and staying in a range. If we had a speaker system on our desk direct to the trading pit, then this is how the noise level would fluctuate in relation to the price action. So what is this proxy that we are using? It is simply the Bollinger Band Width related to the 5-minute time fractal of the price. We will come on to Bollinger Band Width in Chapter 10 but suffice to say here that it is one of the best measures we have of volatility as expressed by range contraction and range expansion (not by the traditional method of historical volatility).



Source: Bloomberg L.P.

Figure 3.16 US Long Bond Future and Pit Decibel Meter (Bollinger Band Width)

The point of all this is that volatility, as measured by the expansion and contraction in the range of market prices, captures the movement of price from trends to ranges and so forth. Moreover, it is capturing the human animal spirit that creates these volatility cycles, in the sense that it is the crowd behaviour and the group psychology that is constantly at play. It is the basic human animal nature that creates cycles in market volatility as the fluctuations between periods of high activity and low activity are driven by the crowd and the crowd behavioural dynamics.

However, before we move on to look in depth at how we can best identify the probability of this volatility, it is worth outlining the psychology of trading regimes.

THE PSYCHOLOGY OF A RANGE-TRADING REGIME

When a market trades in a range the participants that make up that market are in a mass psychological state of uncertainty, doubt and nervousness. This is not to say that everyone involved in the market will exhibit these neurological traits but, as a

whole, the market will exhibit them. In a range-trading regime, supply and demand in the market are pretty evenly matched and therefore for every person who wants to sell and does sell, there will be another person who wants to buy and does buy. More accurately for every dollar, pound, euro or yen that belongs to the sellers there will be a relatively even amount of dollars, pounds, euros or yen that belong to the buyers. That is to say, it is not the *amount* of buyers and sellers that have to be roughly evenly matched, it is the *power* in terms of capital that has to be roughly evenly matched. The market could have 100 people lined up as sellers and their total selling power comes to \$100m, but if the market has just one buyer and her total buying power comes to \$100m, then guess what? Supply and demand is matched and the market doesn't go anywhere. Nonetheless, even though not everyone in a range-trading regime will have an underlying uncertain psychology, a lot of people will and their thought process will be one where a large element of doubt and/or neutrality over the market direction will be prevalent. Either they don't really have any strong convictions about the market and therefore are not prepared to take any positions, or they don't care about the market. This could be because they have just enjoyed a big trending move in the market and have booked their profits, so for the time being they are enjoying the fact that they have money in the bank. One reason a market will pause and range trade when it is in the middle of a big trend is because people become fearful of their profits turning into losses and therefore cash in their chips by closing their position. The people who have done that will not really care about the market for a while because they are feeling very pleased and smug that




Figure 3.17 Range Trading Regime

they have taken a large profit from the market. There will, of course, still be many people in the market and a lot of their psychology will be one of uncertainty. Perhaps they will take a long position but will be quick to close it out when they see a profit because they are not really sure that the market will continue higher. The same goes for the short sellers. They take a short position but will also be quick to close it out when they see a quick profit. In this way the range-trading regime develops. Uncertainty feeds on uncertainty and once people start to see a range, that will fuel even more uncertainty that the market will break higher or lower. The psychology becomes self-reinforcing in a sort of feedback loop, as shown in Figure 3.17.

THE PSYCHOLOGY OF A TRENDING REGIME

When a market is in a trending regime the participants involved in the market are in a mass psychological state of greed, fear, denial, euphoria and panic. When a market trends it is because of a mismatch in the underlying supply and demand for that market. Again, this is not to say that it is the physical number of participants that are mismatched but that the market power or capital being traded is mismatched. Let's say, at a particular moment in the market, there are 100 buyers whose combined capital comes to \$100m but there is one seller whose capital comes to \$200m. Guess what? The market goes down. This mismatch of supply and demand (in this case more supply than demand) will cause market makers to mark down their prices.*

When other market participants see the market going down they could do one of three things. They could ignore it altogether. They could go against the market direction and buy because that is what their thought process is telling them to do. Or they could go with the market direction and sell. If they choose the latter it might be because that is what their original thought process was telling them to do. It is more likely, however, that they had some original thought in their mind to sell, but did not have the conviction to put the position on and act in the market. Now, however, when they see the price going down they have consonant cognitions in their minds that they had wanted to sell, and now they see the price is going down. Time to sell buddy. So they sell. It is also likely that some market participants were neutral, but once they see the price moving they decide to sell because they don't want to miss the move. This is the classic fear of missing the move. There is nothing that grates on the psyche of

* One of the most misunderstood market sayings comes when people explain, usually in a jovial fashion, that the reason a market moved up, for instance, is that there were "more buyers than sellers". This is usually said in a jovial fashion because the common understanding is that buyers and sellers have to be evenly matched and, therefore, there can never ever be "more buyers than sellers". The response is normally frowned upon by economists and other fundamentalists who need a "reason" why the price has moved up but, as with other market phrases, if they stopped and thought about it they would realise that the answer is not as daft as they think it is. What should be said of course is "more demand, than supply" and then it would make sense, surely even to economists!

a trader or fund manager more than missing a big move in the market because being on those moves after all is what they are, or should be, paid to do. I have been in that psychological place many times and I can tell you from personal experience that it is not a pleasant place to be. This is one of the most persuasive reasons I can think of to look at systematic or semi-discretionary trend following as a trading style because it is the only trading style I know that *guarantees* you to be on the correct side and participate in a big trending move in the market. So, by adopting a trend-following approach you can kill stone dead one of the most destructive psychological urges in the markets, which is to look at a move happening and all the time thinking “I should be on this move” only to finally succumb and act at just the wrong moment when the trend ends. Figure 3.18 shows the feedback mechanism for a trending regime.




Figure 3.18 Trending Regime

THE MARKET PSYCHOLOGY CYCLE

To make the market move the way they do, no matter what the timescale or fractal involved, the market participants will go through similar emotions. Day traders will go through this emotional cycle almost every day whereas really long-term investors will probably only go through one complete cycle in their lifetime! This is the “fractalised” nature of markets where the longer term the investment horizon,

the longer term will be the emotions. Long-term investors couldn't care less about day-to-day movements, whereas day traders don't care whether the market is half of what it was 10 years ago.

Figure 3.19 shows the different phases of the psychology cycle and the emotional references are to those being experienced by someone who is assumed to be "long" of the stock market. The same cycle will occur in other markets but the emotional references will differ. It will also obviously occur if someone is short of the market where the emotional references will be reversed. The basis of this psychology cycle is the wave personalities described in the book *Elliott Wave Principle* (Frost and Prechter) but I have added my own bent on the emotions involved based on my own market experience.




Figure 3.19 Market Psychology Cycle

- **Despair**: This is the bottom of the psychology cycle. The psychology of the majority in the market is of impending doom. All the news flow is bad and any good news is totally ignored as the market focuses on the bad news only. The price is in overshooting territory and the psychology is that the market will continue to plummet. It is so bearish that hardly anyone thinks it is even possible for the market to rise. All the so-called fundamentals appear to be the worst.

Characteristic thoughts are: "we're doomed", "this market will never go up again", "how am I going to survive?".

- **Denial/Comfort**: The price has risen and the feeling of impending doom has abated. The market starts to feel more comfortable. However, the psychology of the majority is a state of denial and the expectation is for the market to continue to fall after this rise in the price.

Characteristic thoughts are: “this is only a bounce in a bear market”, “I am not really comfortable although it is a relief”.

- *Happiness*: The price rises some more and the market news flow starts to get a little optimistic. Only a minority of the market is long.

Characteristic thoughts are: “things seem a little better”, “we’re back to the trading range”

- *Greed/Joy*: The price continues to rise as the majority of the market starts to notice the fact that the market has been rising. The psychology is that it becomes imperative to buy the market as the overriding perception is that the market is going to take off. By the end of this stage the majority is long.

Characteristic thoughts are: “if I don’t buy it now I will completely miss the move”, “I really need to buy more”.

- *Euphoria*: This is the top of the psychology cycle. The psychology of the majority is euphoric. All the news flow is good and any bad news is totally ignored as the market focuses on the good news only. The price is in overshooting territory and the psychology is that the market will continue to rise forever (it is virtually impossible for it to fall). All the so-called fundamentals appear to be the best.

Characteristic thoughts are: “we’re made for life”, “this market will never go down again”, “it’s going to the moon”.

- *Denial/Worry*: The price has dropped and the feeling of euphoria has abated. The market starts to worry a little bit. However, the psychology of the majority is a state of denial and the expectation is for the market to continue to rally after the drop in the price.

Characteristic thoughts are: “this is only a retracement in a bull market”, “I am not really worried although it is a bit of a blow”.

- *Sadness*: The price falls some more and the market news flow starts to get a little pessimistic. Only a minority of the market is short.

Characteristic thoughts are: “things seem a little worse”, “we’re back to the trading range”.

- *Fear/Panic*: The price continues to fall as the majority of the market starts to notice the fact that the market has been falling. The psychology is that it becomes imperative to sell the market as the overriding perception is that the market is going to collapse. By the end of this stage the majority is short.

Characteristic thoughts are: “if I don’t sell it now I will be ruined”, “I really need to short more”.

- *Despair*: See previous description!

And off we go again on another psychological cycle in the markets. This cycle in psychology repeats at every time fractal. It is the essence of the ebb and flow of the market price.

Part II

Existing Trading Regime Analysis

*“There are three ways of learning golf: by study, which is the most wearisome,
by imitation which is the most fallacious, and by experience, which is the most bitter.”*

Robert Browning

4

Orthodox Pattern Recognition

The oldest form of technical market analysis (Japanese candlestick analysis) came about as traders started physically recording the price action of rice futures in the 1600s. In the west a similar genesis took place in technical analysis in the 19th century when traders in the stock exchanges started noting the price action in their heads (ticker tape readers) and then noting the price action down on paper with a series of dots to denote higher prices and crosses to denote lower prices. This latter development was the start of what we now call Point and Figure charting and it was the development of this method of recording price action that led the way to what we now recognise as modern technical analysis. All these methods have one thing in common in that their goal is to use the changing dynamics of the price action to analyse the underlying psychology of the market as a whole to see where and how it is changing. This basic element of technical market analysis underlies everything in the field from Gann to quantitative analysis.

In the 1940s two men wrote what is rightly considered to be one of the “bibles” of technical market analysis. Robert Edwards and John Magee’s *Technical Analysis of Stock Trends* put in print for the first time the techniques of analysing market psychology via modern (at the time) price bar charts. A full bar chart contains the four items of information that are important to a student of price action: the opening price, the high price, the low price and the closing price. The three most important prices are the high, low and close, with the open being not as important in bar chart analysis (whereas in candlestick analysis it is vital). One bar will represent one time frame or fractal: a month, a week, a day, an hour and so on (see Figure 4.1).

A time series of the humble bar chart contains a huge amount of information for the technical market analyst and, although not vital to orthodox pattern recognition, it is on this form of chart that Edwards and Magee concentrated, specifically focusing on price patterns that are formed by a series of individual bars.

The idea of price patterns refers to a time series of price that forms in such a way to resemble certain patterns that, in the theory of technical analysis, is the manifestation




Figure 4.1 The bar chart

of market psychology and therefore can be used to anticipate future psychology and prices. As we shall see as we move through the subsequent chapters, there is a very real overlap between orthodox pattern recognition, Elliott Wave analysis and Dow Theory in that, for instance, consolidation phases in Elliott are identifiable patterns in orthodox analysis. A triangle in Elliott is a triangle in Edwards and Magee terms (albeit with one major difference with regard to symmetrical triangles) but Edwards and Magee's original work took more inspiration from Dow's original work rather than from Elliott's brilliant addition to it.

Patterns in the time series of a market price are quite simply the market coming to life. It is human nature to need a reference point, and so when we think of the price of something then we instantly compare it to previous prices. "Have you seen the price of fish recently?" "I remember when a pint of beer was half the price it is now!" "My word the price of laptops has fallen a long way." If all we had was the current price and no memory of past prices then how would we make a judgement on whether this price is where we should strike a deal?

However, the best way to help us to understand about the psychology behind price patterns is to look at some examples, so let's look at some of these orthodox price patterns and see how they can help us to identify the type of trading regime we are in or are likely to be in.

TRIANGLES

Triangles are the quintessential range-trading regimes. They occur when the market psychology becomes so uncertain that a triangular pattern in the price occurs as the price contracts in ever-decreasing ranges. There are three types of triangular patterns that the technical or trading regime analyst should be aware of.

Ascending Triangles

Ascending triangles generally occur when the market has been in an uptrend (or experienced at least a partial up move before the pattern) and they are defined by flat topside resistance and rising uptrend support.

The psychology behind an ascending triangle is generally bullish because the resistance (selling) on the topside of the range is being conducted around the same level, whereas the support (buying) on the bottom of the range is being conducted at an ever-increasing level. Essentially, the market is being capped by net selling around a static level but when the market falls it is being met with net demand at ever-increasing price levels. This is a bullish sign and normally means that the price breaks out of the topside of an ascending triangle.

Although the final directional resolution of an ascending triangle can help us, as trading regime analysts we are really interested in the fact that price action is coiling up and what that means for the future regime. When we see a coiling in price action like this it usually means that the current trading regime is a range and the next trading regime will be a trend. After price range contraction comes price range expansion. So the question is how can we anticipate the change into a ranging regime then out of the range and into a probable trending regime.




Figure 4.2 EURJPY monthly – classic ascending triangle

Figure 4.2 shows a classical ascending triangle in the EURJPY exchange rate on the monthly time fractal (each bar represents one month). Notice that the bars on these charts do not show the opening level but only the high, low and close. For the purposes of orthodox pattern recognition this is fine, but when we come to candlestick analysis the opening level becomes important. We can see that the market peaked in the middle of 2003 just above 140 before declining to around 125

towards the end of that year. In 2004 a rally to look again at the 140 level, and then a fall which stopped short of the 125 level was the confirmation that we, as trading regime analysts, would need to make us increase the probability to above 50% that we are in a long-term (because we are using monthly time fractal charts) range-trading environment.

The market did indeed trade in an ever-decreasing range for another two years before finally breaking up out of the triangle in 2006. So, in this example, the trading regime analyst would have concluded at point P that there was a higher than average probability that EURJPY was in a range-trading regime and that, as a consequence, risk management and capital allocations to trading styles should be adjusted accordingly. When the market broke up out of the triangle the probabilities switched to a trending regime and the analyst could switch risk management techniques again. Anyone already familiar with orthodox pattern recognition will know that a study of volume can add a lot to the analysis, especially at breakout zones, and I will cover some of these issues later. Suffice to say now that when a breakout point is reached, and the market breaks out of the pattern on increasing volume, it adds confidence to the fact that the breakout is a real one.

Now, a word about support and resistance lines. In my opinion, the mistake a lot of people make when drawing support and resistance (or trend) lines on price charts is that they get too exact with them. For example, the exact top in the EURJPY example (Figure 4.2) in 2003 was 140.96. When the market declined then rallied again in 2004, the rally stopped at 139.04. So, at this point, the analyst has a decision to make. Does he draw a descending line from 140.96 through the 139.04 top or does he conclude that, in the context of a *monthly* fractal, 139.04 is pretty close to 140.96 so it *could* be a horizontal line (and therefore a potential ascending triangle). The answer, as always with orthodox pattern recognition, has to be pragmatic. In this example the analyst would have held out the possibility of a symmetrical triangle (see below) developing, but that potential would have been negated in late 2004 when the price rallied back up to the 140.96 level only to fall back again. At that point the analyst would be close to 100% certain that it was an ascending rather than a symmetrical triangle that was in play in EURJPY case.

So I like to think of support and resistance levels in terms of price *zones* rather than specific price points. Doing this means that I can deal with the reality of what happens during market trading. Again, let us take the EURJPY example. In late 2004 when the market rallied back up to the 140.96 level it actually rallied just beyond this level to an intra-month high of 141.61. The way markets work in practice is that people, human beings, put on a trade (they take a position either long or short the market) and as a corollary to that they would generally have a stop-loss order. A stop-loss order is, as the name suggests, an order that “stops” the “loss” on your position. There are a multitude of various ways that you can structure your “stops” but having a plan to stop your losses is absolutely essential to long-term survival in the trading or investment business. Every financial disaster in history has come about ultimately because someone somewhere took their eye off the ball and did not have

the discipline to control their losses. In terms of actually actioning stop-loss orders, this is where it gets interesting because a lot of people will leave what are called “standing orders” in the market. That is they will instruct their broker to action the stop-loss order if the market hits a certain level. This is especially true in the foreign exchange market because of the 24-hour nature of the market – the market is open five and a half days a week from the opening in Wellington New Zealand on Monday (Sunday evening in London) to the close in New York on Friday night (Saturday in New Zealand) – so a big move can happen when people are away from their desks, at home, asleep. As they do not want to wake up the next morning having lost a huge amount of money, they will leave stop-loss orders overnight to limit the loss.

Now, the banks or brokers with whom people have left these standing stop-loss orders will obviously know where these orders are located and, generally, people leave stop orders at levels where they think: if the market goes to this point, it should go further. For example, if a person is short of the market he will generally leave a stop-loss order above the previous significant high, thinking that if the high is beaten then the market should continue to rally well beyond it. Therefore, the stops are generally clustered around certain levels. The main point here is that the banks will know where the stops are, so in a quiet market when there is not a great deal of volume going through, the banks (the market makers) will usually try to move the price towards the levels where the stops are located. I used to think this was a conspiracy theory made up by people who had been stopped out many times before but, having seen it in practice, I can assure the reader that this does happen. Then the stops are “filled”, there is a price spike as that happens but then, when the market realises that there are no follow through orders, the market drifts back to its previous level. This is what is known in the foreign exchange markets as “stop hunting” and the people that do it, the market makers, are known as “sharks” because, for them, it is a feeding frenzy of profit.

Think about it. The banks know there are standing orders to buy EURJPY, for example, 30 pips above where the market is trading and it is a quiet market. They, the market makers, slowly drift the price up to that level via their bid offer quotes. Hey presto, the customers buy orders get filled (the market maker has sold, gone short, EURJPY at that level) but, when all the buy orders are “done” the market realises that there are no more buyers in the market and the market then plummets back down. The market makers, who have gone short EURJPY by filling their customers buy orders at those lofty levels, have made a lot of money. Cha ching!! I can tell you from personal experience that it is not a pleasant experience for the person who gets his stop-loss order filled only for the market to reverse back in the direction of his original position. This, as many of us will know, can be one of the most frustrating aspects of trading! The thing is, however, that everyone knows it goes on yet people still persist in leaving standing orders at obvious price levels. However, this is a bit of a “catch 22” for market participants because the consequences of not having a stop-loss can be far, far worse than getting stopped and then the market reversing. What a lot of people do is to have “mental” stops where they do not leave orders with

banks (or brokers in stock markets) but have a level in their mind that, if the price goes there, they will stop the position out themselves. This is fine when you are in the office, but what do you do when you are asleep?

So getting back to our EURJPY example, undoubtedly the market spiked through the 141 level in late 2004 as stop-loss orders were triggered then, when the market realised that no one else was actually buying EURJPY, it fell back down. What does the analyst conclude? Well, by being pragmatic and keeping an open mind the analyst can take account of these price spikes. In this example, the resistance level of 140.96 should not be thought of as a specific price point but more of a price *zone*. The market has gone up past the level but, given that it came back down below it pretty quickly, we can conclude that, in all probability, this price spike was due to stop-loss orders being triggered and, therefore, the resistance “zone” is still intact. Analysing the psychology of the market via price pattern recognition is, unfortunately for the more mathematically minded, much more of an art than a science.

One way I like to think about these support and resistance zones is by putting myself in the position of being a technical analyst in the days before computers. In those good old days before software applications had tools to allow drawing lines on price charts (not until the 1980s did things begin to take off in this area) poor old technical analysts, or chartists as some were known, would be plotting charts and drawing lines on those charts on a hard copy of graph paper. I remember an old boss of mine in the 1990s who would still diligently plot prices of the UK Gilt market on his graph paper every day. He had been doing this for years and, as a consequence, he had this amazing record of UK Gilt prices filled in by pencil on a series of hard copy graph paper that, because of the scale and time series factors, had to be stuck together with tape from one bit of paper to another. As a result, this self-made chart was about 10 feet long when it was laid out end to end! But it was an irreplaceable record of price action and, in fact, a lot of people even now say that physically plotting price action in this way is a much better way of getting to “*feel*” the price action and therefore the *psychology* of the market for yourself. I have tried doing this from time to time and I agree that it gives the analyst a very good feeling for where the psychological support and resistance zones are, but the obvious problem is that it is not very efficient. For day traders concentrating on just one market it can be a very useful, albeit incredibly intense, discipline to use just as traders on the old stock exchange floors used to adopt this method via their intra-day Point and Figure charts.

By putting myself in the position of being a technical analyst in the days before computers, I can think of how the support and resistance lines would be drawn by using a pencil and a ruler rather than a piece of software. The reason I think about support and resistance zones rather than specific price points is that I assume I am simply drawing these lines with a *thick, blunt* pencil! If you were drawing lines on a chart with a very fine pencil, then the problem is that you can become too specific and not take into account price spikes from stop-loss orders being filled just below or above previous lows or highs. In the EURJPY example in Figure 4.2, drawing a

horizontal line from the high of 140.96 with a “thick pencil” gives us the pragmatic outcome of allowing for a little spike above the specific level before deciding that the resistance zone has been broken.

The point is that the analyst has always to maintain an open mind as to the potential outcomes and adjust his analysis on the basis of an updated price action. I know this is anathema to our quantitative analyst friends but it is an unavoidable consequence of orthodox pattern recognition analysis. Most technical analysis is not meant to be a purely scientific exercise. Due to the fact that the market is driven by the decisions of human beings, when we are analysing the market we are dealing with and analysing a living, breathing beast. Human psychology is what we are analysing and human psychology is in academics, like economics, a “social” science. Therefore, a certain amount of discretion has to be put in place when analysing anything about the market place. Yes, certain quantitative techniques can and do add a whole lot of value in trading models and analysis (mainly because they take the human emotional element out of running risk by the way) but their downside is a failure to judge certain price movements, for example, a price spike that takes out stop-loss positions as being just that.

As an aside, the interesting aspect I have found about quantitative models is that there is still a human element needed in defining model parameters and defining risk limits among other things. The classic story that comes to mind is the Long Term Capital Management debacle in 1998 when they ran their mean-reverting type models with such high risk in illiquid markets that eventually the mark to market losses forced them to close the positions causing quite a major panic about the banking system in the USA. It is my belief that running a quantitative model is analogous to driving a car. The car can be the best engineered Ferrari ever made but if the driver is an arrogant idiot who thinks he is indestructible and drives well above the speed limit, then it is undoubtedly a recipe for disaster! Someone, a human being, still has to drive the quantitative model and act on risk control procedures (like slowing down or hitting the brakes of a car once in a while!).

Figure 4.3 is another example of a classic ascending triangle, this time in the weekly time fractal of the US listed stock Procter & Gamble. We can see that during 2001 and 2002 the market had been in a strong trending regime moving from a low of around \$28 to \$47 where the market peaked in June 2002. The stock then fell back to around \$36 before rallying back up to the \$47 zone in August and September 2002. It is here, at point P–, that the analyst should notice that the \$47 zone is acting as a strong resistance to the price and, therefore, given the previous move down to \$36, the probabilities of the market price action and psychology developing into a contracting ascending triangle have risen. The market subsequently moves down over the next few months until it reverses the down move around \$39 in March 2003 at point P. When the market rallies from point P, the analyst should now be very confident that this market is in the midst of an ascending triangle and, indeed, the stock rallies back to \$47 where, yet again, that zone acts as resistance and the market continues to range trade.




Figure 4.3 Procter & Gamble weekly – classic ascending triangle

At point P– we have then an inkling that the market might be setting itself up for (as this is the weekly fractal) a number of months of range trading and by point P we are pretty sure that this is the case. So, in this example, risk management and capital allocations to trend-following or range-trading styles could have been adjusted accordingly from point P– onwards.

When we get to the breakout point the probabilities switch round dramatically from a range-trading regime to a trending regime and we can see that over the next few months Procter & Gamble moved in a trending fashion from \$47 to \$57, some 21%.

Descending Triangles

Descending triangles generally occur when the market has been in a downtrend (or experienced at least a partial down move before the pattern) and they are defined by flat support and falling downtrend resistance.

The psychology behind a descending triangle is generally bearish because the support (buying) on the bottom of the range is being conducted around the same level, whereas the resistance (selling) on the top of the range is being conducted at an ever-decreasing level. Essentially, the market is being supported by net buying around a static level but when the market rises it is being met with net supply at ever-decreasing price levels. This is a bearish sign and normally means that the price breaks down out of the bottom side of an ascending triangle.

Although the final directional resolution of a descending triangle can help us, as trading regime analysts we are really interested in the fact that price action is coiling

up and what that means for the future regime. The question is, therefore, how can we anticipate the change into a ranging regime and then out of the range into a probable trending regime?




Figure 4.4 USDJPY weekly – classic descending triangle

Figure 4.4 shows a classic descending triangle in the USDJPY exchange rate in the weekly time fractal. We can see that the market was in a strong downtrend throughout 2002 until the low around 116 in June of that year. Then the market reversed up to the 125 level, before falling to 120 and then rallying back up towards 125 at point P- around December 2002. However, the market stalled here and this would have given the trading regime analyst a clue that, after such a strong downtrend that ended mid-year, the market could now be in an extended range-trading regime. When the market then falls to the 116 area and finds support there at point P around March 2003, the probabilities of USDJPY being in a descending triangular range trading regime increase markedly to well above 50%. Of course, by then the market has already been in a range-trading regime for a number of months, so should we not seriously question the longevity of such a regime extending into the future? Again, the pragmatic nature of analysing and trading the markets comes to the fore. This is not an exact science but an exercise in probability scenario analysis and identification and so, really, from point P at least the trading regime analyst should be altering his probability spectrum to reflect the fact that the regime is in the process of switching. Even from point P, though, when the analyst should be very confident of the fact that a range-trading regime exists, there is still money to be made by trading the range and not allocating too much capital, if any, to trend trading because the market does not finally break out of the descending triangle until September 2003. Yes, we had a bit of a wobble around May 2003 when the market fell briefly below 116 but it quickly recovered and went back into the trading range, meaning that the move down

was probably a price spike to take out stops (see discussion above about support and resistance zones, rather than specific points).

Symmetrical Triangles

Symmetrical triangles generally occur after the market has been in trending period either up or down (or the market has experienced at least a partial up or down move prior to the pattern developing) and they are defined by ever-decreasing levels of support and resistance, leading to a classic coiling of the price.

The net psychology behind a symmetrical triangle is one of complete uncertainty because the support (buying) on the bottom of the range is being conducted at ever-increasing levels whereas the resistance (selling) on the top of the range is being conducted at an ever-decreasing level. Essentially, the market is coiling up by a process of net buying around an increasing level when the market falls and net selling at ever-decreasing price levels when the market rises. This is a sign that the battle between the bulls and the bears is very evenly matched and, because of this neutrality, it normally means that the price can break out of either side of a symmetrical triangle.

There is actually considerable debate regarding symmetrical triangles in the technical analysis community, not concerning their obvious uncertain psychology nor about their resolution in a trending period, but about the direction of the breakout. Some technical analysts, specifically those who follow the Elliott Wave Principle, regard symmetrical triangles as continuation patterns that will resolve in a trending move in the direction of the move in which the market entered the triangle. So if the market was in an uptrend before the symmetrical triangle then the triangle should be resolved with a breakout to another uptrending move. In my market observations, and there are many examples to draw on, what can be identified in Edwards and Magee style as a symmetrical triangle can sometimes break out the topside and sometimes break out the bottom side, and this reality I think reflects the psychology behind the symmetrical. What the symmetrical triangle is telling the market analyst is that “we, the participants in the market, are extremely uncertain as to the future course of the market direction”. Moreover they are saying “we are so uncertain that there are not enough buyers to break previous peaks and not enough sellers to break previous troughs”. The market is so wound up on uncertainty that there is no statement of intent such as with ascending or descending triangles that show a bias for demand and supply. Therefore, when the break comes it can be either way.

A good example of this is can be seen in Figure 4.5, which shows a symmetrical triangle in the EURUSD exchange rate from 2000 to 2002 using the weekly time fractal.

We can see that the market was in a long downtrend throughout 1999 and marks a low in October 2000 around 0.83. The market then rallies up towards the 0.95 level before retreating back down to just below 0.85 in June 2001 at point P—. At this point




Figure 4.5 EURUSD weekly – classic symmetrical triangle

the trading regime analyst should be alert to the fact that a triangular range-trading regime could well be in the making and should already be adjusting his risk parameters accordingly. Then, when the market rallies again this time only up to point P around 0.93 and stops there, the trading regime analyst is fairly certain that this whole price action has been and indeed still is a symmetrical triangle. This means two things to the analyst. Firstly, because we are still in the triangular price pattern, the odds favour a continuation of the range-trading regime. Secondly, however, it also means that when the symmetrical triangle is over, the odds dramatically shift in favour of a trending regime. So while there is still some money to be made from trading the range, the analyst knows that at some point the market is going to break out into a trend. Point P comes in around October 2001 and the breakout occurs the following summer, so in this example there are still a number of months of range trading to be taken advantage of. When the market finally breaks out of the range-trading regime it immediately enters a strong trending regime, which lasts for a number of weeks before another range-trading regime takes hold.

I remember this pattern very well because, as it was forming, I debated quite strongly with some advocates of Elliott Wave, who were labelling the price action as a symmetrical triangle when it was in progress, that the symmetrical need not necessarily lead to a continuation of the long-term downtrend and could just as well break upwards. We shall talk more about volume considerations later but an analysis of volume can give you a really good clue in symmetrical triangles as to which direction the ultimate break will occur. The rule of thumb in volume analysis is very logical. If volume increases during a price move then the conclusion should be that the direction of prices is where the balance of power lies. That is, volume increasing

on up moves and decreasing on down moves would suggest to the analyst that the real underlying market forces are such that demand is stronger than supply. During this time in the EURUSD exchange rate, the volume increased during the up legs of the triangle and down during the down legs of the triangle, giving the analyst a big clue that the probable resolution of the triangular pattern would be to the upside. It is worth noting that although the Elliott Wave analysts may have labelled the price pattern as a symmetrical triangle when it was forming, when it broke to the upside it would have been re-labelled to reflect that reality. However, far from being a weakness of Elliott Wave as many suggest, this “alternative count” aspect is, as we shall see later, one of its attributes.




Figure 4.6 Gold weekly – classic symmetrical triangle

Figure 4.6 shows another example of a classical symmetrical triangle, this time in the gold market. Gold is a contentious subject these days as most people tend to believe that it has no value whatsoever in modern life. The gold detractors say the yellow metal is a relic that pays no interest and has no industrial value like copper or even silver. On the other hand, there is a small but vocal minority who believe that gold was, is and always will be the ultimate store of value in the world and that it holds an undeniably emotional attachment to human beings that has not changed for millennia. They say that the gold is the ultimate money and that the fiat (by government decree or government backed) system of money that our generation has grown up with is only “worth” something because of the confidence that people have in it. If that confidence was to disappear, the gold bugs say, then the system would be shown for what it is – a system of government IOUs without any backing of physical, hard assets. Most people dismiss these gold fanatics as cranks, crackpots and conspiracy theorists but I like to keep an open mind and, while realising that the probability of the global economic system returning to some form of gold standard

(where the currencies are backed by gold) is low, the possibility is always there. One thing is for sure: the price of gold (not just in US dollar terms but also priced in a basket of currencies) tends to fluctuate with the *liquidity* (as measured very broadly by money supply and credit) in the world, making the gold price a valuable indicator of economic conditions and, therefore, worth watching. It is said that the great “maestro” himself, the ex-Chairman of the Federal Reserve Alan Greenspan, watched the price of gold to give him a sense of liquidity and credit conditions, and the Fed Funds rate under his chairmanship moved broadly coincidentally with the gold price. Whether Greenspan did enough to contain the bubbles in stocks, housing and credit of course is another question, and may, in the years to come, return to haunt the man they called the maestro.

Returning to our gold price chart in Figure 4.6 we can see that the market was in an uptrend in 2004 and into 2005 before topping out around 460 dollars an ounce. The market then falls to around 420 in February 2005 before rallying back up to 450 at point P—. At this point the trading regime analyst should be thinking that the possibility of a triangular price pattern developing is getting higher but at this stage it remains only a possibility. The market then drifts down towards the previous low of 420 but does not come below it, stopping around 423. This is point P in June 2005 and at this point the analyst sees that we have the makings of a symmetrical triangle. He should therefore be aware of the fact that range trading will probably continue for a while, but once it ends there should be a period of trending regime. Subsequent to point P, the gold price rises and falls within the contracting bounds of the symmetrical triangle until in August 2005 it finally breaks out of the pattern and into a strong trending period. Notice that, subsequent to the initial break, the price moved back initially but that the previous upper bound of the triangle then held as support. We will examine this phenomenon of old resistance becoming support (and vice versa) later, but suffice to say now that when this happens after a breakout of a price pattern it is extra evidence that the market is about to enter a strong trending period amid increasing volatility.

This example shows again that the difficulty in identifying triangular price patterns is that the hard evidence of the pattern comes quite a while after the market has been range trading and does not give the trading regime analyst much more time to benefit from the range-trading regime. However, perhaps the main benefit for the volatility analyst is that the identification of the price pattern provides superb priming for the inevitable breakout into a trending regime.

RECTANGLES

The price action of a market sometimes goes into a prolonged period of sideways action where the upper and lower bounds are horizontal or very close to horizontal. For that obvious reason such patterns became known as rectangles. The psychology behind a rectangle is one where buyers and sellers are in competition (as always) but

that competition is very fierce indeed. That is, the power of the buyers and the power of the sellers does not seem to dissipate as the range trading goes on, as is the case with the converging ranges of the triangular patterns. In triangles the market is in an increasing state of uncertainty and lack of confidence, whereas in rectangles the equilibrium price zone is formed by a psychology of very confident sellers battling head to head with very confident buyers.

This psychological battle is so fierce that when it is resolved it usually means that the losing side (either the bulls or the bears) capitulate to such an extent that the market goes into a strong trending regime. Therefore the identification of rectangular price patterns is useful for the trading regime analyst in spotting range-trading regimes and trending regimes.




Figure 4.7 NZDJPY daily – classic rectangle

Figure 4.7 shows the bar chart of the New Zealand dollar exchange rate with the Japanese yen in the daily time fractal and the vertical scale shows how many yen it takes to buy one New Zealand dollar. Rectangles are usually continuation patterns (a pause in the longer term trend) and so they will usually be preceded by a trending regime. In Figure 4.7 we can see that the New Zealand dollar (the kiwi or “the bird” in the jargon of the foreign exchange market) was in a solid uptrend (it was taking more yen to buy a kiwi) in March 2007 when in late April it trades around the 89 level. The chart shows that there were three days when the price could not extend beyond the 89 level, and so at this point (P–) the trading regime analyst should start to think about the possibility of a range developing. Then, on the third day, the market drops to around 87.50 but on the next day the market falls through the 87.50 level before quickly recovering to close above 88. Over the next two days the market

has changed from range trading to something new, probably a trending regime. We can see that the share price of Tesco held around this new 200 support level for a number of weeks in 2000, even though it dipped briefly below a few times before recovering quickly.

As has been mentioned, the price action or technical analyst should never get too hung up on a precise level of support or resistance because he is analysing a dynamic, living animal which is a human market place, so, rather than the precise numbers themselves, it is more the general *form* of price action that is important because that identifies the overall psychology of the market. In this regard we can see that when the market dipped back into the old range it generally did so on an intra-week basis with the weekly close being back above the zone. The closing price is, as we shall see later, the most important price for the technical analyst and so by the middle of 2000 the trading regime analyst will have realised that the old rectangular trading range regime is over, with the old top of that rectangle now acting as support for the market. This places the odds for a trending regime well above the odds for another range-trading regime so the analyst can adjust his risk parameters and game plan accordingly. As can be seen, the share price of Tesco subsequently enters a period of a strong upward trend giving the trader or investment manager a great opportunity to benefit from that trending regime.

FLAGS

A flag is another range-trading regime that comes after a strong trending move, and that normally leads to a continuation of the trend.

Although it is worth noting that it can have a slight slope to the pattern, a flag price pattern is essentially a scaled down rectangle in that it occurs over a short period of time in relation to a rectangular price pattern, and the reason for this is the context in which it appears. When a market is in a particularly strong trend it can have a tendency to exhibit a strong move followed by a short pause and then a strong move again. The market psychology is so strong that the pause in the trend is short lived and when the pause occurs it merely pushes the market sideways for a short time before resuming the strong trend. Therefore, one of the most important elements in identifying a flag consolidation pattern is that the preceding move must be a strong, almost vertical move. Like other areas of pattern recognition, a flag can sometimes be identified when there has not been this preceding strong move and the analyst must watch for this because in that case the underlying psychology of the price pattern will be something else entirely. The old saying is that “a flag must fly on a mast” and so there must be a preceding strong move before the pause if we, as trading regime analysts, are going to be confident that the subsequent move out of the pause will be another strong trending regime.

Figure 4.9 shows another bar chart of the New Zealand dollar exchange rate with the Japanese yen in the daily time fractal during July and August 2007. We can see

at point P—. At this point the trading regime analyst should be putting quite a reasonable probability that a rectangle is forming because the market has been capped a couple of times at the same level on the upside and has been supported a few times on the downside at the same level. In January 1999 the market moves up again towards the 200 level and yet again it holds as resistance at point P. It is at this point that the trading regime analyst will be very confident that a rectangle is in force and he should be adjusting his risk parameters and game plan accordingly. Of course, the trading regime analyst does not know how long the rectangular range-trading regime could last as all he knows is that the current trading regime is obviously a trading range, but just identifying that significant probability allows him to factor it into his overall risk management game plan. As it turns out, the rectangular range-trading regime goes on for over a year from point P, giving the trader or investment manager plenty of time to take advantage of the consolidation pattern.

In March 2000 the share price of Tesco starts to break above the 200 level for the first time in a couple of years and this gives the trading regime analyst the first clue that a change from a range-trading regime to a trending regime is under way. However, the market does not immediately roar higher from the break above 200 and this, as it turns out, gives the trading regime analyst even more ammunition to suspect that a strong trend is about to develop. Very often a market will break a support or resistance level and then, not so long afterwards, the price will drift back down to that old level again. If we think about what is happening in the market, then it makes sense that we often have this “re-test” of the old support or resistance level. The market breaks higher above 200 and this inevitably brings lots of buyers into the market at or around the same time. So the market initially breaks higher towards the 230 level on this large demand in the market before slowly drifting back down. The drift back down is almost inevitable too, because once the initial demand surge on the break of 200 has been satiated it leaves a shortage of buyers relative to sellers, or a shortage of demand relative to supply, in the market. When that happens there is only one thing the market can do and that is fall back, no matter what the rationalising fundamentalists or anyone else is saying. So the market falls back but *crucially* it stops falling around the 200 initial breakout level. When a market stops falling we know it is because there is demand coming into the market to meet, and maybe exceed, the supply of that market, and if that demand is coming in around a previous breakout level then it is a very, very strong clue that the market is well supported and therefore has the potential to move into a strong upward trend.

What is actually happening here is that the people who stood by and watched as the market broke higher from 200 to 230 (the people that would have been saying to themselves “damn, if I had just bought it on the breakout but I didn’t, if it comes back down towards 200 then I am going to buy some”) are coming in to buy the market. For them, their psychology has changed from one of uncertainty or neglect when the market was trading in a range, to one of certainty and commitment when the market has broken out of that range. So, crucially, old resistance has now become new support and that gives the trading regime analyst a huge clue that psychology

again tests the upper limit of 89 and fails to break through before it falls to the 87.50 level, which now holds as support. This is point P and it is at this point that the trading regime analyst should be quite confident that a rectangular price pattern is developing because we have resistance around 89 and support around 87 to 87.50. So, from late April the trading regime analyst will be confident that “kiwi/yen” is trading in a range with well-defined support and resistance zones. He can then look to trade that range using valid discretionary range-trading techniques or by “turning up the risk on his range-trading quantitative models” until the evidence suggests that it has been broken and a new regime is unfolding.

The history shows that the market traded sideways between 89 and 87 to 87.50 for over a month before, in early June, breaking up decisively. This is the breakout level and, at this point, the trading regime analyst should be switching the probability from a range-trading regime to a trending regime. The market has had its long consolidation and is now ready for the next impulsive move higher. The market subsequently moves from 89 in early June to 97 by mid-July, a move of 10% in a little over a month – a huge move for an exchange rate.




Figure 4.8 Tesco weekly – classic rectangle

Figure 4.8 shows a classical rectangular range-trading regime in the share price of the UK-based supermarket chain, Tesco, from 1998 to 2000 using the weekly time fractal. Notice how the market was in a strong uptrend going into the consolidation phase but then the peak around 200 (British pence) in early 1998 was not bettered when it was re-tested a number of weeks later. That would have been the first clue that this market had the possibility of entering a range-trading regime. The market then trades down to around a low of 150, which holds as support a few times in late 1998




Figure 4.9 NZDJPY daily – classic flag

that during the early part of July the exchange rate was trading in a range between around 94 and 98. Then on 26 July the market broke down dramatically with some follow through over the next couple of days to give a big 8.5% down move in the space of just six days. In the context of exchange rates this is a very big move in such a short space of time, so at this stage the trading regime analyst, knowing the price action elements that make up a flag, should be on alert for a potential flag pattern developing. The market made a low on 30 July and over the next few days started to find support around that low of 89.27 and resistance around a high of around 92.40. This is classic flag price action when the market has made a large, fast move and then the market psychology switches to a horizontal and narrow (in the context of the preceding move) range-trading regime. On 9 August the exchange rate closed below recent range closes and this was the first warning that the market was breaking down again. It continued down again the next day and the next before accelerating dramatically down to produce a massive 15% down move in the space of, again, just six days.

The psychology of this price action is fast and furious, which is why flags (and pennants) usually occur in the shorter time fractals and not really in the longer ones. The way I like to think of flags and pennants is that they are analogous to a pit stop in Formula One motor racing. The driver has been racing the car at 200 miles per hour before entering the pit lane, getting the mechanic team to change the tyres and refuel, and then racing out again straight back up to 200 miles per hour. In market terms the crowd psychology takes a big initial move (possibly a very rapid short covering move or in stock markets possibly in reaction to an earnings warning or some exogenous shock) and it does so on an increase in volume. Then when that initial supply or demand has been satiated, the volume drops quite quickly to

make supply and demand much more balanced. However, the people who wanted to buy (in the case of an initial up move) or sell (in the case of an initial down move), but were not quick enough to jump on board, are feeling very nervous at this point because they want the market to retrace to allow them good levels to enter the market. However, the psychological pressure to “get into the trade” is so intense that they don’t wait too long before going to the market and pulling the trigger to open their positions (or indeed for others to close existing ones). There is another surge in volume as the flag breaks and the market has another fast and furious move in the direction in which the market entered the flag.

The obvious question is: How does the trading regime analyst take advantage of this price pattern? Due to the fact that these patterns usually occur in the shorter time fractals, then only the very short-term trading regime traders will be able to take advantage of the range-trading regime that manifests itself in the flag pattern. If the analyst notices a particularly violent move in the daily time fractal, then he should be on “flag watch” and after a couple of days of range trading should be fairly confident of the range top and bottom parameters. This then gives the short-term (probably day) trader some good risk and reward entry and exit levels from which to try to get some value out of the range while it lasts. More importantly, though, it gives the short-term trader a huge opportunity to benefit from the subsequent strong trending regime and upsurge in volatility after the flag is completed.

Pennants

A pennant is similar to a flag in that it is a very dynamic structure of market psychology and will therefore generally occur in the shorter time fractals. Again, it can be regarded as a scaled down symmetrical triangle although, as with the flag, it can have a slight slope (like wedges, which we discuss later).

Figure 4.10 shows the share price of Northern Rock, the UK bank, in the daily time fractal from March to October 2007. We can see that the market had started to fall at an ever-accelerating rate from March and by the low on 16 August had fallen by some 45% in five months. In fact, the low on 16 August marked a fall of 25% over the previous six days and so, with these two falls put together, the market had met the first condition of a pennant – namely, that the move preceding the pattern should be very impulsive. The market then traded up for a few days, down for a few days and then up again to trace out a fast, sideways movement in price that resembles a pennant. The market then broke down and utterly collapsed a matter of days after that with two huge gaps separating one day’s trading session. The share price fell 83% in 14 days leaving the market stunned with how a major UK bank could have collapsed so quickly.

However, the trading regime analyst would have noticed the pennant formation in the middle of this collapse and, as the market started to break down from the pennant formation, should have been anticipating another major trending move. And it is




Figure 4.10 Northern Rock daily – classic pennant

some trending move – well, for the very short-term traders it is! As mentioned, the characteristic of flags and pennants is that they are usually confined to the shorter time fractals so they reside in the realm of the day traders and short-term swing traders. However, due to the dynamic nature of the volatility breakout they can be very profitable patterns to identify.

Another important aspect to point out with regard to this collapse in Northern Rock is that, contrary to most people's opinion, this should not have come as a major surprise. The one thing that continues to astound me about markets is that when so-called surprises happen they do not turn out to have been such a surprise after all. A lot of people refer to the tail events, fat tails or long tails. These refer to the tails or sides of the famous bell-shaped normal probability distribution and are the rare events that seem to come out of the blue. Nassim Taleb in his book *The Black Swan* refers to these events as "black swan" events, referring to the discovery of black swans in Australia in the 17th century when, until then, the Western world had assumed with certainty that all swans were white. The contention is that predicting these events is impossible because they essentially are new discoveries.

My contention, however, is that these tail events in a financial or commodity market context do *not* come out of the blue to anyone who was following the market price action. The *trend* of the stock price of companies that blow up and go bust is almost always down for a long time before anything hits the fan.

Figure 4.11 shows the weekly time fractal bar chart of the share price of Enron, the US utilities company that went so spectacularly bankrupt in the early years of the 21st century after it emerged that some of the company's senior employees were engaging in a massive fraud. The chart also shows the 5- and 40-week moving averages of the share price and we can see that they provide a very good indicator of the direction of

fundamental stock analysts turn round and say to themselves that “you know I can’t work it out, the numbers that the company have in terms of corporate profitability and balance sheet are all far too good to justify the share price being down this far. I wonder what is going on? Why is the stock being continually sold?” Perhaps the most relevant aspect to come out of the whole sordid Enron affair should be that at what point should stock analysts turn round and ask themselves that question?

Look, the share price of Enron was in free fall for a long, long time before the news broke that the company had been misleading everybody with their accounting practices and so, when the company was telling everybody that everything was fine, the actual holders of the stock were bailing out and selling (or people were borrowing stock in order to short it). This is the only way the share price could go down. It’s called supply and demand. A stock gets sold by investors and traders for many, many diverse reasons but, in total, if supply is greater than demand we can say that the collective opinion of the market is that the stock is not a good investment. So why should this fact be somehow ignored by these analysts? Are they all correct and the sellers wrong? I suppose they might have been but the *fact* remains that the stock was being sold. One of the reasons that people look at the price action as a lead on the fundamentals is because of potential insider selling. That is, if someone in or close to the company smells a rat and/or knows the game is about to be up soon, then it is not totally unheard of (but of course it is illegal) for them to offload their stock before the bad news starts to come out. Now, I am not saying that this was happening in this case, but it might well have been. The fact is that we do not know who is selling and for what reason but the one thing we can say for certainty is that they would not be selling it if they thought it was a great investment. Forget about corporate ratios, balance sheets and credit ratings. There is only ever one number and one number alone that you can have absolute faith in that reflects, at that moment, the truthful outlook for the company. That number is the share price.

The other aspect about this Enron chart is the clear head and shoulders pattern that is visible and a great sign of a clear topping process. We shall examine the head and shoulder in detail later but it is worth noting here that not only did the neckline of the head and shoulders pattern break in March 2001, but the market then went on to have classic re-test of the break zone in May of that year before coming off hard. This break of a previous support zone and then a re-test to hold as resistance, as we have discussed earlier, is a great indicator of a coming strong trending regime, and so it proved here.

The bottom line with Enron is that the bad news was “in the price” when it eventually came out. It is the stock market’s job to discount the future and in the case of Enron it did a great job. It started discounting the fact that the company was going bust for months before the news came out. If you were a trend follower then you would have been in cash or short of this stock for many months before the end came. The big long tail events or black swans or whatever you want to call them do *not* come as surprises to people who are following the trend of the market. In fact this is where trend following makes its money. The strategy relies on extended trends or the fat tails to make extraordinary profits in order to make up for the small losses




Figure 4.11 Enron weekly – the only number you can trust

the long-term trend. The 5-week moving average is the thinner line and the 40-week moving average is the thicker, smoother line. When the 5 is above the 40, the trend pressure is considered to be up as demand is outstripping supply, and when the 5 is below the 40 the trend pressure is considered to be down as supply is outstripping demand. There are times when the 5 slips below or above the 40 only briefly before reversing, but when there is an extended trend then they do a very good job of showing the path of least resistance of the market. We can see that in March 2001 the 5-week moving average slipped below the 40-week moving average and indicated the beginning of the end for Enron. The trend, as measured by the demand and supply dynamics of these moving averages, remained down until the share price reached zero. So the long-term trend of Enron was pointed firmly down for *months* before the story finally broke that the books had been cooked in November 2001.

What is truly jaw dropping is that as late as 26 October 2001 the *Wall Street Journal* ran a story saying that of the 17 analysts on Wall Street who followed the stock, 10 had a “strong buy” or equivalent rating, five rated it a buy and only one had a sell signal! Presumably the other one was neutral on the stock. The majority of analysts who covered the stock had also remained with a “strong buy” or “buy” for the preceding months as the share price was tumbling. However, before we criticise the analysts too much it must be pointed out that the numbers they were working with in terms of corporate profitability and balance sheet came from the company itself and, of course, these turned out to be fraudulent. What else had the analyst to go on? Even if they met with the executives on a daily basis they would be given the same old falsified numbers. So to these analysts who analyse things on a fundamental basis the stock looked ever-greater “value” as the share price continued to tumble. However, this is where my sympathy with the stock analysts disappears. At what point would

it occurs on the majority of its trades. Therefore, if trend following exists to take advantage of extreme situations, does it not make sense for market analysts to at least be cognisant of strong trends developing? Surely if they did that then they would at least be aware, in an Enron or Northern Rock type situation, that there was, in fact, very good reasons why the share price was tanking down.

BROADENING PATTERNS

The broadening pattern or expanding triangle as it is sometimes known is quite a rare beast and the reason for this rarity is that the psychology behind it is quite uncommon. However, when it appears it can provide for some high probability situations that allow the trading regime analyst to benefit from such identification.

The psychology behind an expanding pattern is quite interesting because it usually occurs at long-term turning points in the market, be that at whatever time fractal the pattern occurs. That is, if it occurs on the hourly time fractal then it would normally establish a high or low for many, many hours to come and, similarly, if it occurs on the weekly time fractal then it would normally establish a high or low for many, many weeks to come. This is because the psychology behind the price action is one of either euphoria in relation to a top or despair in relation to a bottom. In other words, the emotions behind the expanding pattern are extremely powerful and, in this sense, the price action represents the culmination of something big.

Figure 4.12 shows the bar chart of the share price of Microsoft (MSFT) in the weekly time fractal from 1996 to 2001. We can see that the market was in a long-term uptrend in the years working up to 1999. In fact the share price of Microsoft



Figure 4.12 Microsoft weekly – expanding top

was around \$2.50 at the start of 1993 and so, by January 1999 when it reached \$44, it had risen some 1,760% in only six years. It was, as we all know, the story of the decade and the darling of the internet bubble generation. At the start of 1999 the real acceleration of the internet bubble had yet to occur but MSFT and its founder Bill Gates were the iconic images of the internet explosion that was occurring. These are the kind of socioeconomic situations, the popular social and cultural developments, where major tops (or bottoms) can occur. Just because a company or market reaches the kind of iconic status that Microsoft did at the end of the 1990s and has had a major long-term move, doesn't mean there has to be a major turning point near, but it is a clear sign to a market analyst that the emotion and psychology invested in the "story" is, in all probability, near an extreme. In the 1940s, Edwards and Magee described the characteristic of a broadening formation as a market being "out of control" and "a situation in which the public is excitedly committed and which is being whipped around by wild rumours". In fact what a broadening formation is tracing out is the same psychology that produces larger price bar ranges at the end of a trend. That is, as the trend becomes mature, realised volatility will sometimes tend to accelerate (from already high levels) as prices fluctuate wildly in a mass battle between the bulls and the bears. We can also see this in other indicators, such as Average True Range which we will discuss later.

We can see from the chart that Microsoft reached a high of around \$44 in January 1999 before falling back to \$36 and rising to just short of \$48 in April. From there it fell back to around \$38 and held. At this point the analyst can draw two lines: one connecting the January and April highs and one connecting the subsequent reaction lows. The noticeable thing about these lines would be that they diverge from one another in an expanding fashion and this would be the first clue for the trading regime analyst that a longer term broadening formation was developing. So at this point the analyst should be telling themselves two things. Firstly, that the potential broadening formation should provide some good opportunities to trade the range and, secondly, that the eventual break of the broadening formation should provide for quite a strong trending regime. Of course the astute technical analyst will expect and anticipate that the break will come to the downside in the case of a broadening top and to the upside in the case of a broadening bottom, because these formations in psychology represent the culmination of a trend. Incidentally, Edwards and Magee point out that broadening bottoms do not occur in stocks because the psychology behind the pattern is suited only to tops (i.e. a very excited public getting involved). I can appreciate the point they are making and it is quite unusual to find broadening formations at the bottom in individual stocks; however, *relative* markets like currencies can still exhibit this psychology because a top in one currency is going to be a bottom in another. Equally, relative plays in other markets can exhibit the same psychology.

Getting back to our Microsoft example, from the lows that held in May 1999 the market moved up to over \$50 in July, back down again to test the lower support line in August, up to \$49 in September, back to the lower line in November at \$42 and then zoomed up to a fraction short of \$60 for the end of the century, which

also happened to coincide with the resistance of the upper broadening line. This top would turn out to be the top in Microsoft, which at the time of writing in 2008 has yet to be neared again yet alone beaten. It was a historic top in all senses and one that probably will turn out to be hugely significant in the future in the sense that if and when it is beaten, it will be a huge clue that some sort of mega bull market might well be returning. As mentioned, broadening tops and bottoms exhibit such strong emotions and psychology that they tend to be the price limit for a long, long time.

From the high in MSFT at the broadening top extreme the share price fell back down quite swiftly, but again held at the lower line, allowing another opportunity for the trading regime analyst to benefit from the trading range, but this time the bounce failed to beat the previous top and the market retreated. On the week of 7 April 2000 Microsoft closed below the lower line of the broadening formation and this should have meant that the trading regime analyst was seeing flashing red lights all over the place. The fact that the market closed below the lower boundary meant that there was a probability that the break of the broadening formation was occurring.

At the opening bell the public tell, but the close is for the pros. What this means is that the closing price action of the session is the most important part of the day because the closing price is the price that people go home and think about or fund managers factor into their valuation models. At the opening bell there is usually a flurry of orders to get filled and a great many of these orders will have come from the public who, having done their analysis the night before, give buy or sell orders to their brokers to be done at the opening the next day. In this sense they are telling the market at the outset what their opinion of the market is, but like anything else in the markets it is the physical supply *of* and demand *for* the traded asset that is of the paramount importance. Suppose there was a flurry of sell orders in a stock at the opening bell to be filled. These are filled at lower prices than the previous close because the brokers know that there is supply coming and so will mark down the prices. When all this initial supply is out of the way, however, and it turns out that there is no more coming through then guess what the price will do? It has to rise according to the laws of supply and demand. The closing price, on the other hand, is the one that people think about when they go home at night and therefore has a much deeper psychological impact than the opening.

So once Microsoft closed below the lower support of the broadening formation, the odds on a strong trending down move had massively increased, and so it turned out to be the case with the market moving down from \$44 to \$30 (or some 32%) in only seven weeks.

HEAD AND SHOULDERS

A huge amount has been written and analysed about the venerable head and shoulders price pattern over the decades, which makes it probably the most famous of all orthodox pattern recognition price patterns. A lot of what has been written is, for

want of a better phrase, complete and utter rubbish, written by fundamentalists who have an agenda to discredit technical analysis as much as possible but, at the same time, a lot of users of technical analysis talk about this statesman of a pattern without really knowing what they are talking about. As a result there is quite a bit of confusion over this pattern, which is a shame because it has led to the important and classical psychological attributes of this price behaviour becoming lost in the mire.

Quite simply the head and shoulders pattern is the classical manifestation of trend exhaustion and reversal psychology. It doesn't always, but quite often, lead to a significant top or bottom in a market and therefore is usually followed by a significant period of trending behaviour as the psychology of the market changes from bull to bear or vice versa. In actual fact the period of price action when the head and shoulders pattern is forming is more often than not a range-trading regime and this makes sense when we consider what is happening to psychology and market behaviour during this time.

Let's get one thing straight. The head and shoulders price pattern is a *reversal* pattern. That is, it will not always, but often does, lead to a significant reversal in the previous significant trend and, therefore, like all reversal patterns, it *HAS* to have something to reverse. I get very frustrated when people in the markets talk about head and shoulders appearing in the price action when there has been no previous significant trend, or worse still when they point out a head and shoulders pattern coming after a significant downtrend. This is complete nonsense and displays a clear lack of understanding that the head and shoulders pattern represents the psychological top or bottom of *previous trends*. The head and shoulders pattern must have something to reverse and if the previous trend has been sideways or down then what on earth is it supposed to be reversing? If it was a reverse head and shoulders pattern coming at the bottom of a downtrend then fair enough, but I have had people try to point out to me a normal head and shoulders pattern at the bottom of a *downtrend* with the conclusion that the price will fall. This lack of understanding regarding the context in which the pattern must appear has also led to studies on the viability of the head and shoulders (mainly by economists) giving misleading negative results. I cannot stress this strongly enough – a head and shoulders price pattern has to have something to reverse.

If we assume that a market has experienced a significant uptrend and that the market is making a series of higher highs and higher lows (a normal definition of an uptrend), then the market makes a higher high but reverses off it still to a higher low, but perhaps on an increase in traded volume. The market then tries to recover perhaps on lower volume, but fails to reach the previous high. If the market then reverses on another increase in volume and breaks down through the previous low, then the market has just been through a classic head and shoulders pattern. The psychology changes from a strongly bullish (the left shoulder) to an element of doubt creeping in (the higher volume on the down leg of the head) then to denial that “everything is still ok” (the up leg of the right shoulder) and finally to outright bearish (the down leg of the right shoulder). The failure to make a new high is a clear sign, but not

definitive, that the market may be about to reverse significantly. A more definitive sign comes with the break of the “neckline” which is formed by drawing a trend line underneath the pattern. Finally the last piece of the jigsaw comes when the market makes a lower low from the previous one.

Figure 4.13 shows the US Dollar Index (an index of US dollar exchange rates versus a basket of currencies) using the weekly time fractal. We can see that the US dollar was in a strong uptrend in the late 1990s and into 2000, but it then peaks in 2001 around 121 on the index. The market then drops to still make a higher low than the low of early 2001 but the subsequent rally fails (just) to make a new high above 121. When the market starts to come off from that peak in early 2002 then the trading regime analyst should be gearing up for a potential trending regime and should be adjusting his risk plan accordingly. The break of the neckline around 113 leads to a strong downtrend period that lasts for a couple of years and gives the trading regime analyst many opportunities to benefit.




Figure 4.13 USD Index weekly – classic head and shoulders

This is a classic example of a head and shoulders price pattern and the psychology just jumps out from the chart. The market is in an uptrend and the bulls are firmly in control, but the market then fails to make that last new high. This is the first sign that the bulls are not as strong as they once were and might potentially be losing control of the market. When the neckline is broken at the breakout level, the odds are very firmly in favour of the bears and, when the price breaks the previous low, the bulls are dead in the water. Notice that, when the head and shoulders pattern is forming, it appears to be trading in a range-trading regime. The market is fluctuating up and down within a range but the range is quite wide in relation to the previous uptrend.

In effect the market is experiencing higher volatility as measured by the range of price movement (not the rate of change of prices) and this is quite a common phenomenon at the end of trends. When a trend is becoming exhausted the price action tends to become a little more volatile in this particular ranging fashion as the people who have been long the market for a while cash in their chips to people who want to buy the market, even though it has moved (in this case up) a great deal already.

Coming at the top of an uptrend this is classical “distribution” when the strong hands (those that have been long for a while) sell to the weak hands (those that are just catching on to the “story” and want in at any price). In stock market terms the strong hands are generally thought of as the institutional players who, at the end of an uptrend, sell (distribute) their shares to the general public who are considered to be the weak hands. This is the basis of the age-old adage that when the public get involved in a stock market rally then it is generally time to sell, and this makes sense when we consider what is happening. In the early stages of a market (it need not just be the stock market) rally, no one is interested but when the market rises a great deal and the press get hold of the story, it has the effect of sucking in marginal players (those that have half an eye on market developments but not full-time market watchers) like the general public who are hearing how much the market has gone up and no doubt hearing bullish predictions for the future. The public then buy the market but those that are selling to them are the distributors, like Lord Rothschild mentioned previously, who have had the strong hands to hold the market for a long time. A very important socioeconomic point to understand is that it is the *prices* (of markets) that make the news not the other way round! We never hear in the media of a major trend in the market until *AFTER* it has been in force for a long time. Of course we don’t. This is the point. The market has to have moved a long way before it *becomes* news! Once this “news” hits the streets the bulk of the trend is usually finished and it’s only a matter of time before a major top forms.

Figure 4.14 is another example of a classic head and shoulders in the US Dollar Index. Again we are looking at a chart of the weekly time fractal but the period is back in the 1980s. Throughout 1983, 1984 and the early part of 1985 the US dollar was in a bull market with the Dollar Index reaching a high of 165 in February 1985. Then the market dropped to a low of 145 in April 1985 before rallying back towards the 155 area. At this point the trading regime analyst is aware that there is a potential head and shoulders price pattern in the making, but the possibility still exists that the market will continue its multi-year bull trend. The low in April was a higher low than the previous low in late 1984, so the market is still making higher lows although the analyst can now draw an upward sloping (neck) line through these lows to act as a potential sell signal. Once the market starts to retreat from the 155 level, however, the trading regime analyst is on full alert for the potential that this price pattern will turn out to be a head and shoulders pattern and that the future trading regime will, in all probability, be a trending regime. The fact that the price action since late 1984 had started to range quite wildly should be another warning sign to the analyst that the previous trend could be showing signs of exhaustion and therefore potential reversal.

markets by these countries in a concerted effort to weaken the US dollar because (a) the US current account deficit had become a very high 3.5% of GDP and (b) a weaker dollar would help the US to emerge from quite a deep recession by helping its export sector. As the statement from the meeting read:

10. These positive economic developments notwithstanding, there are large imbalances in external positions which pose potential problems, and which reflect a wide range of factors. Among these are: the deterioration in its external position which the U.S experienced from its period of very rapid relative growth: the particularly large impact on the U.S. current account of the economic difficulties and the adjustment efforts of some major developing countries; the difficulty of trade access in some markets; and the appreciation of the U.S. dollar. The interaction of these factors – relative growth rates, the debt problems of developing countries, and exchange rate development – has contributed to large, potentially destabilizing external imbalances among major industrial countries. In particular, the United States has a large and growing current account deficit, and Japan, and to a lesser extent Germany, large and growing current account surpluses.

11. The U.S. current account deficit, together with other factors, is now contributing to protectionist pressures which, if not resisted, could lead to mutually destructive retaliation with serious damage to the world economy: world trade would shrink, real growth rates could even turn negative, unemployment would rise still higher, and debt-burdened developing countries would be unable to secure the export earnings they vitally need.

And the intended course of action:

18. The Ministers and Governors agreed that exchange rates should play a role in adjusting external imbalances. In order to do this, exchange rates should better reflect fundamental economic conditions than has been the case. They believe that agreed policy actions must be implemented and reinforced to improve the fundamentals further, and that in view of the present and prospective changes in fundamentals, some further orderly appreciation of the main non-dollar currencies against the dollar is desirable. They stand ready to cooperate more closely to encourage this when to do so would be helpful.

This deal came as a major shock to the foreign exchange market and when trading resumed in Wellington, New Zealand, on Monday 23 September (Sunday night in London and Sunday afternoon in New York) the effect was immediate. We can see from the chart that the US Dollar Index gapped lower from about 140 to 134 and then continued its decline for weeks thereafter without a hint of a pullback until February 1986. After that, though, the market continued its downtrend for another couple of years. This downtrend was one of the strongest downtrends seen in the foreign exchange markets in history and made a lot of people very rich indeed because,




Figure 4.14 USD Index weekly – classic head and shoulders

When the neckline of the head and shoulders was broken in the summer of 1985 the market accelerated down on that particular week and the trading regime analysis was skewed strongly to a trending regime developing. The fact that the analysis is being done on the weekly time fractal, and that the head and shoulders pattern took quite a number of months to play out, suggests that the trending regime in which the market now finds itself will likely last for a long time.

Now, what is really interesting about this analysis is that the head and shoulders neckline broke in the summer of 1985 with the technical analysis conclusion that the US dollar would be entering a sustained period of downtrend. We can see that the market fell over the next few weeks, but rallied a small amount back up to the 145 breakout level. This is the re-test of the breakout level that we have talked about which sometimes occurs after a psychological level has been breached and the re-test is a crucial piece of the puzzle because if the, in this case, old support now acts as resistance then we can say with a very high degree of confidence that the underlying psychology of the market has changed. The chart shows that the 145 level acted as resistance for a couple of weeks in late August and early September 1985 but it is what happened next that really excites the technical analyst. Don't forget that, at this point, the trading regime and technical analysis suggests that there is a very high degree of probability that the US Dollar Index is about to undergo a major trending period and the completed head and shoulders pattern suggests that this trend will be down. This is mid-September 1985.

The 22nd of September that year was a Sunday and, unbeknown to the majority of the foreign exchange market, a deal was being struck in the Plaza Hotel in New York between five countries. This deal entailed joint intervention in the foreign exchange

as the central banks were effectively saying “we are going to weaken the dollar”, it was like writing a blank cheque to market participants who chose to follow them.

However, the most important aspect to all of this for the market analyst is that the Plaza Accord, as it became known, occurred after (*AFTER*) the head and shoulders pattern in the US Dollar Index had been completed and for the students of price action this is one of the all-time great examples of how the price *discounts* what is happening now and potentially in the future. The purpose of a market is to discount what might happen in the future, and it is astonishing how accurate price action can be in anticipating something that turns out to be taken as news, which confirms the previous market action. In this instance the market price action had already shown students who bothered to pay attention to it that a major downtrend had a high probability of occurring. Then the Plaza Accord came along and everyone went “wow, that is brand new information” and sold the hell out of the US dollar, but the identifiable change in market psychology as measured by the price action had already told us that the US dollar was going to go down. The Plaza Accord merely acted as an accelerator of the established direction that the market had determined for the dollar. Yet most people view the fact that the US dollar depreciated so much in 1985 and beyond as having been a direct result of the Plaza Accord. Yes the Plaza Accord undoubtedly helped the US dollar to decline because some of the people that sold the dollar were these governments who had reached the agreement but, to students of price action and psychology, the behaviour of the market preceding the Plaza Accord was discounting the fact that the dollar would have depreciated anyway in this time period.

Can the price action of a market predict the “fundamental” news that is going to affect that market? Does the price lead the fundamentals? I would argue very strongly that it can and does on many occasions, with this being a prime example. When I argue this point with economists or fundamentalists they just do not get it. “There must be a catalyst,” they say. “There must be a reason,” they say. I say yes, the catalyst will become clear and the reasons will become clear – **after** the price action has discounted these things! What I don’t get is that the economists and fundamentalists seem to disregard one of the basic functions and objectives of a market that is to discount or anticipate the future. This is what markets do. As with Enron and Northern Rock discussed earlier, the news always lags the price.

Reverse Head and Shoulders

The reverse or inverse head and shoulder pattern is, as the name suggests, the mirror image of the head and shoulders price pattern. Rather than occur at the end of an uptrend in price, the reverse head and shoulders will occur at the end of a downtrend in price. As discussed above, it is vitally important that the context in which the price pattern appears fits with the psychology of the pattern and, therefore, a trend

magnificently, if bull markets climb a wall of worry then bear markets slide down “a slope of hope”.

So the bear market that developed from 2000 was quite relentlessly savage and by the middle of 2002, when the S&P had fallen by some 50% from its peak, emotions and sentiment were exceptionally bearish. Those bulls who had been counting on the V-shaped recovery were now talking of an extended U-shaped recovery, with some even talking about an L-shaped outcome (i.e. no recovery!). These are the pessimistic sentiment conditions from which a trend exhaustion and eventual reversal can occur, and therefore are ideal conditions for an analysis of the trading regime of the market. The week of 26 July 2002 was an exceptional week because the market spiked down a huge amount to a low of 775, but by the end of the week had recovered to close just above the opening level. This, as we shall see in the next chapter, is a classic hammer bottom in Japanese candlestick terms and not only that but the week saw very heavy volume. As we shall see from the section on volume considerations, both of these facts increase the probability of trend exhaustion and so, from this point, the trading regime analyst would have been on high alert for the downtrend to be ending and for the probabilities of a range-trading regime to increase. We can see that the market did indeed stop going down from that point, and after a four-week rally up to August at 965 the market then reversed and started to slide again. By the week of 11 October 2002 the market was back to the previous lows and, in fact, crucially made a new low of 768 before reversing back by the end of the week. This also occurred on increased volume and so, by this stage, the trading regime analyst would be very confident that the S&P was trading in a range-trading regime in the weekly time fractal. Thereafter the market rallied again up to early December but, again crucially, made a lower high than the peak in August. At this point the analyst is being presented with a market that has made a marginal new low but the subsequent reaction high has not beaten the previous high. While this increases the odds of a reverse head and shoulders occurring, at this stage the market is tracing out a typical bear market pattern of lower lows and lower highs. However, given the fact that the low in October was such a marginal new low, the analyst should be giving a decent probability to the potential for a reverse head and shoulders developing. However, it is the price action of the next four months that will really increase the odds on a reverse head and shoulders coming to fruition.

With all technical market analysis, an analysis of volume can add a huge amount to our examination of market psychology and can greatly enhance confidence levels of certain outcomes. In orthodox pattern recognition, a study of volume is extremely useful and, as we shall see later, it really brings out the emotions of the market. When it comes to head and shoulders and reverse head and shoulder price patterns the study of volume can give some very good clues as to the likely market outcome. Remember that we are talking here about a head and shoulders pattern that has an eventual reversal in the trend and therefore it is a reversal in psychology. So if the market psychology is going to change, we not only see it through the price action of the market, but we also see it through the volume of the market when allied to that price

action. In a head and shoulders pattern the market is in an uptrend, has a correction, makes another rally to a marginal new high, has a correction, makes another rally to a lower high and then the previous trend reverses. Now, if the market is indeed experiencing a reversal in psychology during this process, the volume on the rallies should be dissipating and the volume on the corrections should be showing signs of expansion. That is, the rallies should have less and less volume or power behind them and the corrections should have more and more volume or power behind them. If this is seen to be happening then it greatly increases the chances that a true market reversal is occurring.

In many ways the process of a head and shoulders or a reverse head and shoulders pattern developing can be thought of as one where the control of the market is shifting ever so slowly from either the bulls to the bears in the context of a head and shoulders and from the bears to the bulls in the context of a reverse head and shoulders. As alluded to earlier, it is, in fact, a process of either distribution or accumulation. In any market cycle you have the concept of strong hands and weak hands in reference to the market players, and the strong hands are the ones who have probably been positioned in the market during a trend and who subsequently close their positions by transferring them to the weak hands who have probably watched the trend developing but did nothing and now have the overriding urge to get positioned in the market. After an uptrend the strong hands of the market want to cash in their chips so they sell (distribute) to the weaker hands that buy from them. Therefore, when a market is in a topping process like a head and shoulders pattern, the volume should be seen to be heavier on the down legs when the “strong hands” are selling to the “weak hands” and lighter on the up legs when the “strong hands” are no longer buying. After a downtrend the strong hands of the market want to close their short positions or invest the cash that they have been holding rather than stock, so they buy (accumulate) from the weaker hands that sell to them. When a market is in a bottoming process like a reverse head and shoulders pattern the volume should be seen to be heavier on the up legs when the “strong hands” are buying from the “weak hands” and lighter on the down legs when the “strong hands” are no longer selling. So the head and shoulders price pattern is just one of the manifestations of this accumulation and distribution process that is the driver of markets, and therefore an appreciation of how volume should behave within the pattern is vital to an understanding of the power of the analysis. It should come as no surprise, therefore, that in the studies that I have seen from fundamentalists and economists who seek to discredit the head and shoulders pattern, no account is taken of this volume aspect that is so crucial to the analysis.

Returning to the reverse head and shoulders pattern in the S&P in Figure 4.15, we can look for volume signs that might add to the analysis. In the leg up from the first shoulder, volume in fact dropped, and in the leg down to the head, volume expanded. This is not what we would expect from a reverse head and shoulders pattern and so at the low in October 2002 the analyst would not be thinking that volume was acting in such a way to back up the analysis of a reverse head and shoulders. However, we can

exhaustion and reversal pattern like the reverse head and shoulder will have to appear after a clear downtrend in the price of the market.




Figure 4.15 S&P 500 Index with volume

Figure 4.15 shows a classic example of a reverse head and shoulders pattern. This is the bar chart of the S&P 500 Index in the USA using the weekly time fractal from 2000 until 2004 with its volume histogram. After the “dot.com” bubble burst and the stock market and the economy went into a nosedive from the turn of the millennium, pessimism had grown with this being reflected as a downtrend in the stock market. These were dark days indeed for any fund manager who was “long only” or who was holding onto the notion that the stock market and the economy would rebound quite quickly. I remember all the talk from the fundamentalists of the so-called “V-shaped recovery” in the economy and the stock market where the expectation was that the economy would recover quickly and everything would be back to “normal”. If ever there was a manifestation of denial psychology in the markets, this was it, and it was thrilling as a student of market psychology to witness first hand what was going on. Think about it. People who talk of “V-shaped recoveries” or fast rebounds will generally, not always of course but generally, be in a situation whereby were the recovery to happen in that fashion they would benefit from it or at least be happy about it. So fund managers whose business success is strongly correlated to the performance of the stock market, or investment bankers whose deals will ebb and flow with the markets and the economy, will have a vested interest in the markets and the economy doing well. Another way of describing cognitive dissonance in psychological terms is to say that sometimes people will “cloud their judgement” with what they *hope* will occur and, as Robert Prechter from Elliott Wave International has described so

see that in the up leg from the head low, volume remained relatively high but when the market started to fall into the final shoulder low in March 2003, volume fell. So during the first four months of 2003 the analyst is presented with a market that now has had relatively high volume on the up leg and falling volume on the subsequent fall. Then in the week of 14 March 2003 the S&P fell hard to a low of 788, above the previous low in October, before recovering strongly into the close (another hammer bottom) and volume spiked up. Not only that, but the next week was a big up week and volume spiked up even more. So here, in March 2003 the analyst would have enough evidence to be thinking that both the price and volume analysis was such that the entire price action of the S&P 500 over the previous eight months was a reverse head and shoulders pattern. If the analyst concludes this, then to borrow the famous “spin phrase” from Donald Rumsfeld, the ex US administration official, the known known is that the market has been in a trading range regime and the known unknown is that, when the pattern is completed, the market will move into a trending regime. Of course, the analyst/investor would have been able to take advantage of some of the range-trading regime of the reverse head and shoulders process as the steps involved in the identification process were building, but he would now be fully prepared for the trending regime that was likely to take place. In other words, he will be ready to take full advantage of the regime switch by, for example, increasing risk on trend-following models/methods or by reducing risk on range-trading models/methods.

Once the neckline of the reverse head and shoulders was broken in May 2003 the S&P 500 trended higher over the next few months, making a gain of over 25% up until March 2004.

Therefore, we can see that the head and shoulders price pattern and its counterpart, the reverse head and shoulders, are good indicators of changes in trading regime and volatility conditions.

DOUBLE TOPS AND BOTTOMS

One of the most powerful orthodox chart patterns is the double top or double bottom because it often signals an important trend reversal that will usually be accompanied by a strong trending regime. It is however vital to recognise that the strong trending regime occurs in the context of when it is a proper double top or double bottom! These patterns are, yet again, very misunderstood because a lot of people who have only a passing knowledge about technical analysis will talk about double tops and double bottoms without recognising and pointing out the characteristics that makes these patterns what they are.

Firstly, double tops and double bottoms are reversal patterns and so, in order for a double top or double bottom to occur, there must be something preceding it for the pattern to reverse. That is, there must be a distinct uptrend preceding a double top and a distinct downtrend preceding a double bottom. Like the head and shoulders pattern, I have had people try to point out to me a “double top” or “double bottom”

falling at that point. At this point the analyst might be thinking that there is potential for a double bottom to be put in place in this market but at this stage that is all it is – potential. For a double bottom to be confirmed, we need to see the market rally up past the reaction high (labelled RH on the chart) which is the level it reached after making the first low. Until that high is broken, the market has a lot of potential to carry on in the downtrend and could be in the making of a descending triangle. The reaction high represents the level where supply of the Aussie really overcame any demand for the currency it was experiencing in its bounce from the low, and, as such, represents a level where people who want to be short the Aussie dollar think that it is worth selling at. It represents a psychologically important point for Aussie dollar bears, and so until that level is broken the evidence is not yet strong enough for the analyst to conclude that a double bottom (a solid bottom to the previous downtrend) is in place.

So the Aussie puts in a second low around 0.48, but then just drifts higher for many months without any strong trend to it. Not until April 2002 does the exchange rate start to knock up against the reaction high level of 0.54 and it doesn't break through it right away. It takes its time to break through it (adding to the evidence that it is a very important psychological pressure point) but when it breaks it accelerates to a high of 0.58 before drifting back down again. This is where it gets really exciting for the trading regime analyst because the drift down after the breakup is crucial to the probabilities of how the market might perform in the future. As pointed out in the earlier discussions, when a market breaks up past a strong resistance point or zone (where there is obvious psychological importance attached to it from both demanders and suppliers) then there will probably be an initial surge in demand which will carry the market up to the point where this demand is satiated. When the demand is satiated, the market can really only do one thing (assuming normal levels of supply), and that is fall. But what we are really, really interested in is if the breakup past the psychological resistance zone has actually changed the underlying psychology of the market and the crucial evidence of this will be whether the previous resistance zone now acts as support for the market. If it does not then there is no overwhelming evidence that psychology has changed, but if it does then we can be very excited because it will be clear that psychology has changed from bearish to bullish, and that the bullish sentiment is only in its very early stages.

In the case of the Aussie dollar, we can see that the market drifts back down to the old reaction high of 0.54 and actually dips below it briefly but recovers the level again quickly. As previously mentioned, whether the market dips below a level or shoots through a level is not as important as whether that dip or shoot is over with quickly. That is, support and resistance levels should be thought of as zones where psychologically important forces are at work and not as exact levels from which the price must not go lower or higher. This takes account of how markets work in terms of stop-loss orders and such like. In this case the market dips below 0.54 for the briefest period so there were probably some stop-loss orders below 0.54 that the market took out but, when the spike down came, no more sellers were found and

when the market has been in a distinct sideways pattern. The clue is in the name. Top and bottom . . . *of a trend!*

Secondly, a double top or double bottom only becomes such a pattern when there is a distinct double top or double bottom to the price action. If a market has been in a long-term uptrend and goes up to, say, 20 then falls back down to 15 before rising again to 20 and falling back to 18, could this be considered a double top? We know that the market stalled at the 20 level on what looks like two occasions, but what evidence do we have that the market is not going to rally from the 18 level back up through 20. At this stage the pattern could be turning out to be an ascending triangle or it could just be a pull back in a trend. What we need is hard *evidence* that the uptrend has put in a double top and that hard evidence will come if the market fell below the 15 level, the low that it made after the first test of 20. This is what we call the *reaction* low (a double bottom would have a reaction high) and it is very often missed in people's analysis of what constitutes a double top or double bottom. However, it is vital if the market is to subsequently behave in the trending fashion that can be expected.




Figure 4.16 AUDUSD weekly – classic double bottom

Figure 4.16 shows the bar chart of the Australian dollar exchange rate versus the US dollar in the weekly time fractal. The vertical scale shows the price of Australian dollars in US dollar terms and so we can see that the “Aussie” was in a downtrend in 1999, 2000 and into 2001 as the price was declining from 0.60 (or 60 US cents) to below 0.50 (or 50 US cents). The market put in a low of 0.48 in April 2001 and the market subsequently rallied to a high of 0.54 in August 2001. The price of the Aussie then tumbled again down to the 0.48 previous low in September and stopped

the market, in that situation of demand and supply imbalance, can only do one thing and that is rally. The 0.54 old reaction high is now acting as support for the market so, as trading regime analysts, we can start to conclude that the underlying market psychology has changed from bullish to bearish, that the bullish sentiment is still in its early stages and that means that there is a vastly increased chance that the Aussie dollar will be about to experience a strong, long-term trending move. As it turned out this is exactly what happened with the Aussie moving from the 0.54 level in the middle of 2002 to the 0.75 level by the end of 2003. This is a classic example of how the retest of a previous support or resistance level can act as a very good indicator of a market about to enter a strong trend-trading regime and how that can give the trading regime analyst an edge on the rest of the market.

This example of a double bottom is in the AUDUSD exchange rate and, of course, as currency markets are relative there are two ways of looking at this. Although it is a double bottom in the value of the Australian dollar, if we inverted the chart it is also a double top in the value of the US dollar! However let's now examine a double top in a non-relative market.

Figure 4.17 shows the bar chart of Citigroup, the diversified financial services company, using the weekly time fractal from 2004 until 2008 with the volume histogram in the lower chart. We can see that Citigroup's share price was in a very mild and choppy uptrend from 2004. These are just the sort of weak uptrends that




Figure 4.17 Citigroup weekly with volume – double top

trend followers hate because although the market is making higher highs and higher lows (therefore an uptrend in the strict definition of one) the nature of the trend is very choppy. It is exactly these sorts of trends that trading regime analysis seeks to identify as weak and add value to people like trend followers by enabling them to, for example, reduce their position size to reflect the fact that the regime is more skewed to range trading or weak trends than strong trends. Nevertheless, in the strict definition, Citigroup was in an uptrend from 2004 and in fact had been in an uptrend from 2002 when the share price was \$23. Therefore, when the analyst is examining potential price pattern recognition he would be aware that a broad uptrend had been in place for a long time.

At the end of 2006 the share price rose quite sharply from the \$50 zone to an intra-week high of \$57 during the last week of December before closing at \$55.70. However, the market could not hold these gains and fell back to a low of \$48.05 during the week of 16 March before closing that week at \$49.53. From there the market rallied again but the rally stalled at a high of \$55.55 during the week of 18 May. As the market then started to fall back during the next few weeks, the trading regime analyst should be thinking that the potential of a double top forming around the \$55.5 level has increased. Importantly, though, the analyst, at this stage, still needs further evidence that the psychology has turned from bullish to bearish and to a trending regime. The market still has to break below the reaction low for a double top to be considered in place and confirmed.

On the week of 27 July 2007 the share price of Citigroup did indeed break down below that level and, importantly, it did so on a marked increase in volume. Now, as mentioned previously, price patterns are rarely going to adhere to precise price points of support and resistance but instead will likely show some spikes of undershoot and overshoot. This reflects the reality of market trading and the general zones of support and resistance which are important for a student of market psychology and trading regime. What was interesting here in the Citigroup example is that the previous high closing level was \$55.70 and the market stalled at \$55.55 on the retest of those previous highs. Although the first intra-week high was quite a bit above the second high, the second high stalled at the previous high *closing* level and the astute analyst would have noted this information with interest.

At the end of July the market broke down through the reaction low which was \$48.05 but then, over the next few weeks, the market chopped around between \$45 and intra-week highs of \$50. Again, however, the astute student of price action will have noticed that the closing level of the reaction low week was in actual fact quite a bit above the intra-week low at \$49.53. This is very informative because when the market chopped around during August and September 2007, the weekly closing levels never got above that \$49.53 level. At this point, therefore, the trading regime analyst would have been very excited because of the facts that (a) the breakdown through the reaction low had been done on increasing volume and (b) the reaction low *zone* was obviously now acting as resistance. This price action really confirmed

WEDGES

Wedges are a close cousin of the triangle and the pennant but they have certain psychological characteristics all of their own. Wedges are the bane of the systematic trend follower's and the quantitative analyst's life because they "look" like a trend but are, in actual fact, a range-trading regime that is slightly upward or downward sloping! There are two types of wedges, a rising wedge and a falling wedge.

The Rising Wedge

A rising wedge normally leads to a trending reversal down in price and a falling wedge normally leads to a trending reversal up in price. The price action of the wedge itself is characterised by ever-decreasing ranges but these ever-decreasing ranges occur as the market is moving in an overall slightly upward or downward fashion. The way I like to think of the psychology behind a rising wedge is to use the analogy of a drowning person (not very nice I know but quite apt none the less). If a person is drowning he will struggle to stay above the surface of the water. Every time he struggles to get above the water line it takes more of his strength and saps his energy. He will go under time and time again but summon up enough strength for another go. Eventually the attempts at staying above the water become slower and more lethargic until finally his strength gives up, he is engulfed by the water and sinks like a stone. Perhaps a more appealing analogy is to think of someone walking up a down escalator and never reaching the top. Eventually, his strength will be sapped enough that he is forced to rest and the momentum of the down escalator takes him all the way to the bottom. In a rising wedge the market is fighting for its life to delay the inevitable decline.

Rising wedges occur when it is becoming harder and harder for the market to rally but it does continue to make slight new highs. What we are witnessing in a rising wedge is very strong demand but also extremely strong supply and, crucially, that the suppliers are relentless in their quest to distribute. The demanders on the other hand are very powerful to begin with but their power diminishes over time as the suppliers gain the upper hand. As with broadening formations, wedges usually come around major turning points in the market's life and at the end of long trends. After a long uptrend for instance, the amateur investors (the massed ranks of retail investors normally) will tend to be the most bullish while the professional investors will tend to be the least bullish, at least in private that is. The conspiracy theorists would have us believe that at the end of long uptrends in a stock or a sector, for example, it is the professional investors who are "peddling the story" more than any others in order to find buyers for the investments they want to sell! In other words, one reason why the poor old retail investors always seem to be the last into a trend is that the professional investors are fuelling the flames of the bullish "story" in order to allow them to offload their stock to buyers before it is too late. The cynic in me gives this

that the double-top pattern was in place and that the odds of a strong trending regime were now very high indeed.

The market certainly entered into a strong trend over the next few months with the share price dropping to a low of \$18 in March 2008 – a drop of over 60% in six months – as the cold, hard reality of the credit or liquidity crunch began to be realised by the majority of market participants. Citigroup and other financial shares fell quite precipitously during this time, with established firms like Bear Stearns effectively going bust, and a lot of people exclaimed disbelief that such price movements could occur. As mentioned before, these fat tail events seem to come out of the blue but on closer inspection we can see that our old friend, the trend, was well established before the real damage started to be done. Look at Figure 4.18 of the US Investment Bank Sector using the weekly time fractal and we can see that if we had measured the trend by weekly moving averages, we would have seen that the trend had turned down in the middle of 2007, well before the situation deteriorated so alarmingly. Of course the other interesting feature of this chart is that it is, in itself, a double top, albeit a messy one, where the reaction low acted as a zone resistance rather than a pinpoint resistance on the rebound. So even at a sector level the price action of the investment banks in the US during 2007 was giving big clues of the fallout to come.




Figure 4.18 US Investment Bank Sector weekly

In summary, therefore, double tops and bottoms can provide a very good indicator of a coming trending regime. However, the analyst must guard against calling the pattern too soon and must wait until the reaction high or low is beaten before he can be very confident that psychology has indeed shifted from bullish to bearish, or vice versa.

hypothesis some credibility and it was undoubtedly true in the old days when “pools” operated in the stock market. JP Morgan, James Keene and Jay Gould were just some of the more famous names to benefit from arrangements made with each other to “ramp” stocks up in order to offload them to an unwitting public! This practice is of course now highly illegal in the developed markets and so I am very confident that it no longer goes on in those markets in organised groups. However, this does not mean that professional or institutional investors, who are normally involved in the bullish story and invested from a relatively (to the retail investor) early stage, do not want to offload their investments after a long uptrend. They have an interest in keeping the bullish story going as they attempt to sell their stock. It might not be organised or consciously done, but it happens more often than not that the institutional players are selling (or distributing) after a long uptrend and that the retail punters are buying. It is true to say that this is just a normal and natural part of the market psychology cycle and that institutional investors do not go out of their way to deceive the public nor that retail investors do not blindly buy when the big boys say it is a good thing. Just like the point mentioned earlier about Wall Street existing to sell stocks (IPOs) to the public, we may want to ask ourselves why so many funds are launched and sold to the public *after* long uptrends have already taken place in that particular market sector.

More plausibly, in fact, this behaviour witnessed in the market is just simple psychology at work. Those investors who are long already (the institutions) are talking the bullish story because of their dissonance in “talking their book” and this fuels the story from the retail side. The retail investor wants a piece of the action and this adds to the frenzy in ramping up the price. However, the institutions can’t buy any more as they are fully invested! All they can do is hold or sell, and if some of them start to sell then, due to the size relative to the retail size, the selling may persist for a while. Enter the rising wedge.

Figure 4.19 shows a classic example of a rising wedge in the Indian stock market. This is the bar chart of the S&P CNX Nifty, which is a weighted average index for large companies on the National Stock Exchange of India. It consists of 50 companies representing 24 sectors of the economy and is colloquially known as the “nifty fifty”. We can see from the chart that the Indian stock market had been in an uptrend that accelerated higher from the middle of 2007. After the peak in late October, however, the momentum of the uptrend slowed dramatically with the market, making only a marginal new high in December before falling back. When on 19 December the nifty fifty held around the 5670 zone the probabilities of a rising wedge had dramatically increased. The market then rallied to make another marginal high around the 6350 level, which was the exact resistance of the upper wedge line. Here, in early January, the trading regime or volatility analyst in the Indian stock market would have been extremely excited indeed because the odds were now really favouring a resolution of this rising wedge in a sharp move in a strong trending regime to the downside. On 16 January the market broke below the lower support of the wedge and closed there while, on the next day, it moved in a sideways action but still

can see that the Aussie bottomed in September 2001 around 56 (Japanese yen to buy one Aussie dollar) and then entered a long uptrend taking it to 70 by early January 2002. A fall back to the mid-67s and then a rally to make a marginal new high on 25 January meant that a line could be drawn between these two peaks. The market then chopped around for a few weeks hitting a low of 66 on 13 February before rallying and falling again to test those lows of 66 in early March, but never actually falling below them. Instead the market had a few days of down moves occurring during the course of the day but closing the day up near the highs. In Japanese candlestick terms these are hammer bottoms and a good sign that a previous down move is exhausting and potentially reversing.

Now I deliberately chose this rising wedge as an example because it does not show a perfect situation. In many books the authors tend to show perfect examples of price pattern identification, but I believe that this belies the imperfect reality of the real world where the analyst is often faced with situations where the potential price and psychology pattern do not fit exactly into an analysis of pinpoint accuracy. I cannot repeat this too many times, the reality of market trading means that prices can spike through levels or zones only to snap back very quickly, and it is these spikes or extremely short-term price levels that the analyst has to take into account when identifying a potential psychology pattern. This might sound like “having your cake and eating it” but to me it simply reflects the reality of market price action and liquidity holes. Therefore, when thinking of trend lines I suggest that the analyst should think in terms of these lines being quite thick almost as if they have been drawn by hand with a very thick pencil. Doing this eliminates the unrealistic assumption that support and resistance levels will occur with pinpoint accuracy and embraces the reality that the majority of financial and commodity market trading is a bit messier than that (cue groans from the quantitative analysts!).

So, what this means is that when the Aussie/Yen rate falls to just above 66 in early March it is plausible that the analyst looks at the fact that the closes on those days were towards the top end of the range and therefore a slightly upward sloping line could be drawn from the previous low on 13 February that engulfs the actual lows of early March. The market did not make a new low in early March but instead reverted over the next couple of weeks to a high of 70.94 on 25 March, giving the mean-reverting trader an ideal opportunity to make some money. The highs on 24 and 25 March occurred right on the now extended trend line from the January high, but when the market subsequently dropped and then rallied again a few days later the price spiked up past this trend line, in fact closing above it for two days in a row on 1 and 2 April. However, on 3 April the market came back down into the rising wedge from where it fell again to support lows around 68.20 on 2 May. Does the spike up above the upper line and the close above it present a dilemma to the analyst? Well, yes it does, and frankly the analyst should always be in a dilemma because he should always be thinking about the alternative route the market could take if the preferred option does not turn out to be right. This business is about probabilities, not certainties.



Figure 4.19 Indian Nifty daily – rising wedge

closed below the lower wedge support zone. At this point the odds of a collapse had increased even more and we can see that over the next three days the market crashed from 5910 on 17 January to a low of 4448 on 22 January, a fall of nearly 25% in three days! This is a stock market crash in anyone’s book and so, once again, the notion that stock market crashes come “out of the blue” is refuted by well-established orthodox techniques for evaluating a market’s psychology.

Figure 4.20 shows the bar chart of the Australian dollar exchange rate with the Japanese yen in the daily time fractal from August 2001 until November 2002. We



Figure 4.20 AUDJPY daily – rising wedge

It is absolute common sense for the market analyst, be he fundamentally or technically oriented, to have an alternative strategy for the anticipation of the future if a certain scenario plays out or doesn't. Yet still the investment profession in general puts pressure on analysts to have a clear and concise "certain" forecast of the future. The analysts, colloquially known as "talking heads" on financial media channels, are under pressure from their firms for whom they are acting as marketing devices and from the media presenters themselves to come up with market calls that are "unhedged". That is, they want the analyst to say, with certainty, that the dollar will drop or the bond market will go up. What planet do these people live on? Prediction with any degree of "certainty" is impossible. All we have is anticipation of the future and, if we are anticipating the future but that future doesn't turn out to be true then, surely the best analysis would be to have an alternative scenario (a plan B if you like). So, in this example, when the Aussie/Yen exchange rate spikes up past the upper resistance line and closes above it, the analyst should be thinking about plan B which would be that this price action is not in fact showing a rising wedge psychology and that a price breakout might be about to occur. The analyst thinks about it, he might even act on it. However, when the price fails to stay above the line and falls back into the range the analyst should be very, very quick to change the probabilities back to Plan A, which is the rising wedge scenario. The failure to break out to the topside is actually an extremely strong signal that the rising wedge analysis is very much the preferred option now. The facts have changed and the price or technical analysts, like Keynes the economic analyst, are not only perfectly allowed to, but must, change their mind.

Technical analysts are market chameleons in the sense that they can change their anticipation of market direction very quickly. They do this because they are looking at the market in an objective manner by examining the mass psychology of the market as manifested by the price action. However, to many people changing "views" like this is seen as a weakness and means that the analysis is in some way flimsy. I have heard it many times over the years from fundamental analysts or economists when the technical analysis presents a plan A and plan B depending on market developments. "You can't have it both ways," they say. But this is exactly the point that, sadly they miss. We must have it both ways. Otherwise how are we going to control the risk and have a game plan if the market goes against us? We must be able to look at one market outcome as a preferred choice and another market outcome as the lesser preferred choice. Then if the preferred choice does not materialise then the lesser preferred jumps up in the probability scale and now becomes the preferred outcome. It is all to do with probabilities. All the time.

So by 4 April the analyst in Figure 4.20 will be quite confident that the rising wedge pattern is still in force and we can see from the chart that there was another couple of mean-reverting "coils" in the price action before breaking down at the end of June. Now, on 24 June the market spiked down during the day but closed back above the (thinly drawn) support line. The next day it closed below it, so just as in the previous topside break when the market closed above the line in early April,

the analyst at this stage should be thinking that the probabilities of the rising wedge ending are increased. Should he still be open to the possibility that the breakdown will not happen? Of course he should be. This is common sense. However, this is where an appreciation of various forms of technical market analysis can come into play because when we look at Elliott Wave analysis (which we will do partially later) it would tell us that a rising wedge, or ending diagonal in the language of Elliott, formation should exhibit five moves up from the low, but crucially that these five-wave moves will be very weak. What this means is that a pattern of price action that exhibits five moves up in an overlapping (wave) fashion and looks like a rising wedge, will dramatically increase the odds that the breakdown is about to occur.

On 26 June the Australian dollar/Japanese yen exchange rate did break down dramatically, falling from around 69.50 to a low 66.90 (a 3.7% move which is enormous for a one-day move on an exchange rate). It then fell in a trend-like fashion over the next month, eventually hitting a low of 62.32 on 24 July (a 10.3% move). If you had been a short-term trader you could have taken advantage of the relentless chop of the rising wedge and then crucially switched to your trending method of trading when the breakout was occurring and taken advantage of the swift downtrend in prices. A trader with just one discretionary method, or who is running a certain style of mechanical model, could have benefited by this appreciation of trading regime by switching his style or altering his position size accordingly.

The Falling Wedge

The falling wedge is just the opposite of the rising wedge in terms of price action, but the psychology behind it has its own character. When the market descends in a falling wedge it is the demanders of the market that win out eventually against the powerful but ultimately limited suppliers. The analogy I like to think about in a falling wedge is one of Rocky Balboa, the fictional boxer from the Rocky series of films, when he is going through one of those long spells of being punched repeatedly as he leans back against the ropes of the ring only to eventually find some energy and come out fighting to overcome his opponent. The demanders of the market are the underdogs like Rocky and the suppliers are his big, bad opponents.

Figure 4.21 shows the bar chart of the share price of Caremark Corporation, the company that operates a chain of drugstores located throughout the United States, in the weekly time fractal from 2000 to 2005. We can see that the price was in a long bear market from the top of \$31 in 2000 and after it bounced from a low around \$11 in late 2001, the price began to decline again from around April 2002. At this stage the analyst would not know whether this was going to turn out to be a falling wedge or not, but he would have noticed that the market was falling in accordance with the trend line drawn off the high zones of April to June. When the market crashed to a low of \$11.78 in July 2002, bounced up to the downtrend line at around \$16, then nose-dived again to \$11.51 in October where a low formed, the analyst could now

the breakout can occur in either direction, the psychology of the wedge that happens against the trend is, like the pennant, ultimately in favour of the previous trend. This gives the trading regime analyst a bit of an edge when it comes to identifying a potential breakout from the consolidation regime into a trending regime. In the case of a falling wedge after an up move, the odds are skewed in favour of an ultimate topside break rather than one to the downside, and so the analyst can factor this into his assessment.

Figure 4.22 shows the bar chart of Cisco Systems share price in the weekly time fractal from mid-2001 until early 2004. We can see from the chart that Cisco bottomed in late 2002 and made a fairly impulsive move up to just short of \$16 by the end of November. There then followed a distinct falling wedge pattern lasting many weeks that, by May 2003, would have prepared the trading regime analyst for a volatility breakout, trending regime move to the upside. By early 2004 the share price of Cisco Systems had almost doubled from the falling wedge breakout level.



Figure 4.22 Cisco Systems weekly – falling wedge

Figure 4.23 shows the bar chart of the Mexican peso exchange rate versus the US dollar in the monthly time fractal from 1995 to 2004. The scale shows how many Mexican pesos it takes to buy a US dollar; therefore, when the rate is going up, the Mexican peso is getting weaker versus the dollar. This is an example of a falling wedge occurring as a consolidation period in a bigger trend and we can see from the uptrend in the chart from 1995 to the top of 10.78 in January 1999 that the Mexican peso was declining in value. Astute chartists will realise that the price action in January 1999 was in fact an “outside month”, meaning that the high was higher than the previous month’s high and the low was lower than the previous month’s low. Not only that, but it was an outside month of the previous four months so this was a very big clue that the previous uptrend was in danger of, at the very least,



Figure 4.21 Caremark Corporation weekly – falling wedge

see that the October low was only marginally below the previous low and so a trend line drawn connecting the two lows and, in conjunction with the upper line, would start to form what was beginning to look like a wedge formation. At this stage the analyst could be quite confident that the trading regime was a range-trading regime and he could advise or trade accordingly. We can see that the market did, in fact, continue to chop around as the falling wedge continued to form with the bottom support zone being tested in January 2003 and the upper zone tested again in March 2003.

However, the market then fell to a low of \$11.25 in April and reversed back up without having made a new low or at least a run at the previous low of January. This was a big clue to the trading regime analyst that the market was about to stop its range-trading regime and enter a period of trending behaviour. Another clue that the falling wedge was about to end came in the fact that the market had moved down in five moves from the top in early 2002 to the low of January 2003. The market had done five downs so from January 2003 the analyst would have been on the lookout for potential breakouts. There was, to be fair, a failed topside breakout in March 2003, but the market moved back inside the wedge the next week allowing the analyst or trader to manage that risk accordingly quickly. However, the real breakout came in May 2003 and when the market did not come back from that initial break the analyst should have been very confident that a period of trending behaviour was about to take place. And so it proved, with the share price of Caremark rising relentlessly over the next couple of years from the breakout level around \$12 to a high of \$31 in August 2005.

Although wedges can occur at major market tops or bottoms (those lasting for a long time) it is quite common to see wedge price and psychology formations as consolidations within a bigger trend of prices. In this respect they resemble an angled triangle or a longer lasting pennant. Whereas, with the symmetrical triangle



Figure 4.23 USDMXN monthly – falling wedge

stalling for a while. At this stage the outside month gives the trading regime analyst a heads up that a range might be on its way and when the market dropped to a low of 9.24 in April 1999 and then to 9.2350 in May 1999, more evidence was given that a downward sloping consolidation phase was becoming more probable. We can see that the Mexican peso traded in a range until March 2000 when the rate made another marginal low at 9.1535, right on the lower trend line that the analyst started to connect before, in April and May 1999. From there the market moved up sharply to a high of 10.1370 in June 2000 before falling back again. Here, in July 2000 the analyst would be very confident that the USDMXN exchange rate was in a very long-term falling wedge price pattern and, while he could be quite confident that this range-trading regime will eventually resolve itself in a break to an uptrend, he must also believe that the period of range trading will continue until the evidence suggests otherwise. When the break eventually came it resulted in a strong uptrending regime in USDMXN over a multi-month period.

SUMMARY

This chapter has looked in depth at some of the most orthodox and common price pattern recognition that occurs in the financial and commodity markets. For the market student or analyst that appreciates the human psychological drivers of markets, these patterns in price bring to life the emotions of the market. In turn, this analysis helps the analyst to determine not only the current trading environment but also whether the price is more likely to trade in a range-trading regime or in a trending regime. In the next chapter we shall examine an even older form of analysis and see that it, too, brings the underlying psychology of the overall market to the fore.

5

Japanese Candlesticks

“He will win who knows when to fight and when not to.”

Sun Tzu

Technical market analysis has its roots in the Japanese rice futures markets of the 16th century, where traders there would pay attention to the price action of the rice future in order to get a clear understanding of the overall market psychology. They knew that it was the *people* involved in the market who were the ultimate arbiters of price, and if their psychology changed then it was a sure bet that this would be reflected in the price of rice. Rather than try to guess about external factors that may or may not affect rice traders' psychology such as crop or rice demand news, they focused on the manifestation of that psychology in the actual price in order to anticipate when and how it would be changing. For example, if there had been a downtrend in the price of rice and one day the market opened around the close of the previous day then fell sharply, extending the trend down, only to close the day back at or near the opening level, then these early technical analysts would conclude that there has been a change in market psychology. The bears were no longer in charge of the market because the low prices that were seen during the session brought in a huge amount of buying, enough to propel the market back up to where the day had started. Sentiment had changed and the bears would now be questioning their bearish views. Students of Japanese candlestick charts will recognise this price action as a classical “hammer bottom” in the jargon of the methodology. One of the aspects I love about Japanese candlestick analysis is the descriptive names that are used to describe the price action. In this case the market has been in a downtrend and has “hammered” out a “bottom”. Harami (meaning pregnant), three black crows, and not forgetting the classic doji are other such names that put a very descriptive element to this unique price action analysis.

Although Japanese candlestick analysis is traditionally used to anticipate trend direction or trend reversal, our purpose as trading regime analysts in looking at

candlesticks is not to forecast or anticipate an up or down price direction as such but is more to anticipate direction or no-direction. As trading regime analysts, what we are interested in is whether it is more probable that the market is in a trending or a non-trending mode, and candlestick analysis can help a great deal in that regard. Japanese candlestick analysis is an analysis of the psychology that is being exhibited by participants in the market, and it is this simple observation that should be able to provide simple yet powerful clues as to the kind of trading regime the market is in, or is likely to be in.

As a starting point ask yourself this question. What happens when you are uncertain about an investment decision? Do you leap in feet first with all your capital and make a big decision? Or do you tend to postpone the decision, do nothing and “sit on your hands”? Generally the answer is the latter. You do nothing. It is, in fact, at these times that those of us from the British Isles remember the age-old northern English phrase of “if in doubt, do nowt (nothing)”. So if you are postponing investment decisions and doing nothing when you are uncertain, do you think others might feel the same way? Of course they do. It is natural human psychology at work. Now imagine that a lot of people in the market felt the same uncertainty. What do you think the market price action would look like? Would it exhibit trend-like qualities where the closing price on the day is far away from the opening price? Or would the closing price be not that far away from the open?

If we can assume that it is exactly this psychology of uncertainty and doubt that leads us to do nothing in our investment decisions, then we can say with a degree of confidence that if this psychology is repeated by many market participants, it is exactly the act of doing nothing that leads to range trading! However this seems like a circular theory. Is it the cart before the horse or the horse before the cart?

Socionomics, which has been mentioned earlier, is a new and exciting take on the “dismal science” of economics and pioneers of socionomics, such as Robert Prechter, believe that it is the social mood aspects of the world that influence the economy rather than the economy influencing the social mood. For a lot of people this makes sense because economics is, after all, a “social” science! So socionomic advocates tend to see things through a back to front aspect compared with a traditional economist. For example, whereas a traditional economist would say “recessions make businessmen cautious”, a socionomist would say “cautious businessmen make recessions”. As trading regime analysts we can use this thought process when it comes to using Japanese candlesticks to determine the trading regime. Does range trading lead market participants to be cautious or do cautious market participants lead to range trading? Do trending markets lead market participants to be bold in their decisions or do market participants making bold decisions lead to trending markets?

The answer lies somewhere in between when it comes to trading regime analysis. Yes a market that is range trading can be exhibiting cautious trading (i.e. low volume) but it also can be exhibiting bolder trading (in the form of higher volume). The key is whether the buying and selling volume is being matched on both sides. This is

where Japanese candlestick analysis becomes a great tool for the trading regime analyst because, although volume analysis is still very useful, the simple nature of candlestick analysis allows us to focus on the overall market psychology.

Figure 5.1 shows the information contained in a Japanese candlestick chart. Each candle is a time fractal (a day, a month or a minute) and each candle shows the open, high, low and close for that particular fractal. When the close is higher than the open, then the real body of the candle is not shaded, and when the close is lower than the open, the real body is shaded. The lines from the real body to the high and the low are the wicks.




Figure 5.1 The candle chart

The information in a candlestick chart is just exactly the same as the information in a bar chart only it is presented in such a way that is really brings out the psychology of the individual time fractal. Bar charts are useful but they do not allow the individual psychology to come to life the way Japanese candlestick charts do, and it is this aspect of bringing the individual fractal psychology to life that makes candlesticks so useful for trading regime analysis.

To illustrate this point let's think of the classic Japanese candlestick pattern called a "doji". Doji in Japanese translation means "at the same time" or "together" and it really encapsulates the meaning of uncertainty. A doji candlestick occurs when the open and close on a time fractal (day, month, minute, etc.) are the same or very similar. If we stop to think about the fact that the open and close are the same or very similar, it becomes obvious that the price action is a manifestation of an uncertain market psychology. When the open and close on a trading session, for example a day, are the same, then the buyers and sellers have met and exchanged views (expressed in actual buying and selling) and the net result is that the closing price is the same as the opening price. Supply has met demand (the participants' views) and balanced almost perfectly. In the language of an economist, this market can be said to be, on that day, in perfect equilibrium.

market produced a classic doji candle with the market opening at 988, trading in a 30-point range and closing at 989.

Now, much like identifying support and resistance levels in orthodox pattern recognition analysis, I believe that the analyst should not get too hung up about specifics when it comes to identifying candle patterns in general and doji patterns in particular. I can hear the quantitative analysts screaming as I write this but, in my opinion, while quantitative techniques for managing money can have distinct benefits, market analysis is much, much more of an art than a science. So when it comes to identifying doji (where the perfect doji occurs when the opening price and the closing price are *exactly* the same), if there is a difference between the opening price and the closing price but that difference is so small as to be insignificant (yes, I realise that the definition of insignificant is a subjective decision), then the analyst should not let that close his mind to the psychology of the price action. And if it is the psychology of the price action that interests us more than a mathematical equation, then, as always, the basic rule of thumb to employ is common sense. If the difference between the opening price and closing price is so small as to be insignificant relative to the trading range of the candle, then the market is shouting “*uncertainty!*” at us. The market is saying that it is uncertain enough to close the day, week or whatever very, very close to the opening level. The market tried going up, the market tried going down but the market could not find any direction. It is uncertain and therefore range bound.

Again, the precise accuracy of what constitutes a doji has to be taken into context with the range and scale of the market. My rule of thumb is: if it looks like a doji on the chart then it’s signalling uncertainty in market psychology and, for the sake of argument, is a doji. So the doji week in Figure 5.2, coming after quite a sustained uptrend, is a sign that the market as a whole has suddenly gone from a psychological state of certainty to a psychological state of uncertainty. This in itself does not automatically mean that the market will then go into a period of range trading but it is most definitely a warning sign to the analyst that market psychology is potentially changing. Put another way, it increases the odds of the trend coming to an end and that, in itself, increases the odds of a period of range trading developing.

Therefore, in the example above, when the market opens the week at 988 and closes the week at 989, given that we traded in a 30-point range that week, this price action constitutes a doji and we can say with confidence that, at the close of the week, the human beings throughout the world who make up the trading community in the S&P 500 equity index are, as a group, very uncertain as to the future path of the market. If the market as a whole is exhibiting a very uncertain psychology, then it is a fair bet to conclude that this uncertainty will continue into the next few trading sessions (or time fractals). In this case the market subsequently trades in a horizontal manner for 12 weeks, giving the trader many opportunities to take advantage of the trading range.

I think one of the common mistakes that people make in looking at a doji pattern is to assume that it always means that a price *reversal* is imminent. All candlestick

It is for this reason that, when they appear, dojis are very interesting. Legend has it that in some Japanese trading rooms when someone spots a doji they mention it to the person sitting next to them by saying, in an extended fashion, “do. . .ji”. That person then mentions it in the same extended manner to the next person and so on until the whole trading floor are very soon on their feet chanting “DO..JI, DO..JI, DO..JI”. Whether we as trading regime analysts want to do the same thing is entirely up to the individual, but what we should definitely do as analysts is pay attention when they occur because the humble doji is quite a reliable tool for the trading regime analyst to have in his toolbox.

What a doji is saying to the technical analyst is that “right now, this market is uncertain as to what direction it wants to go”. Now, the traditional way in which to use this information has been for the analyst to use it as a trend reversal indication. That is, if the market has been in an uptrend with strong up days being generated (as demand exceeds supply), and all of a sudden the market gives out a doji (demand and supply in equilibrium), then the analyst would conclude that because of this new uncertainty there is a good chance that the previous trend will reverse. However, while it is true that the previous trend can reverse quite sharply when a doji occurs, it is also true that the appearance of a doji (uncertainty) can lead to a prolonged consolidation in the market price. So when a doji appears after a previous trend in the price it can lead to a transition from a trending regime to a non-trending or ranging regime.

Figure 5.2 shows the candle chart of the S&P 500 share index in the USA in the weekly time fractal. We can see that the market was enjoying quite a strong uptrend throughout April and May 2003. Then in the week of 13 June the price action of the




Figure 5.2 SPX weekly – doji trading range




Figure 5.4 US 10-Year Yield weekly – doji trading range

enough time for the trading regime analyst to add some value from trading the range rather than looking for trends.

Notice that I have circled the doji that occurred at the bottom of the chart in March 2004. In this instance the doji turned out to signal a *reversal* of the trend from down to up rather than a trading range or pause in the trend behaviour of the market and, as alluded to above, this is always a possibility with doji price patterns. We must always remember that the doji price pattern is signalling uncertainty in the behaviour of market participants. Sometimes that uncertainty is resolved quickly, as was the case here in March 2004, but sometimes, and in my observations more likely, the uncertainty takes a while to resolve itself, meaning that the market will likely trade in some sort of range for a while after a doji appears.

ENGULFING AND OUTSIDE BARS

Another splendid price pattern that was recognised by the Japanese some four centuries ago is the “engulfing pattern”, which refers to the situation where the extremes of the current period’s trading engulfs the previous period’s extremes *after a preceding trend in the price*. I highlight the context of where this pattern must appear (i.e. after a preceding trend in the price) because, like so many price patterns, the psychology that is represented in these patterns is only relevant if one puts it in the context of what came before.

There are two types of engulfing pattern: bullish and bearish. The classic bullish engulfing pattern occurs after a preceding downtrend in which the first bar of the pattern has a close lower than its open. Then the second bar of the pattern has an open which is lower than the previous period’s close and the current period’s close

keep an open mind and be aware that, depending on the market in question, perfect dojis very rarely occur. In any case, the price action in June reinforced the severe uncertainty that this market was exhibiting by closing the month only 2 pips away from the open in a 285-pip range from high to low. A 2-pip difference is equivalent to the bid offer spread in the wholesale foreign exchange markets, so in my book that constitutes a perfect doji.

These three doji months in the USDCAD exchange rate from April through to June in 1994 really showed just how uncertain the participants in the market were about the future direction of the exchange rate, so the trading regime analyst will have noticed this and will be processing this information along with his other regime analysis methods and indicators. At this stage the probabilities of a reversal of the previous trend or a sustained period (remember we are looking at the monthly time fractal) of range trading have gone up a lot and, as it turns out, the market did trade in a flat trading range for a long period of 43 months (over three and half years!).

The observant reader will notice that a doji candle appeared during the earlier uptrend in May 1993. In actual fact the difference between the open and close was 13 pips in a 196-pip trading range but close enough, using the same methodology as above, to possibly be considered a “doji type” price action. As it happened, this did not result in an extended period of range trading or a trend reversal and it should be stressed that doji observations will not always result in such consequences. Remember, all we are doing (in fact all we are doing in every form of market analysis) is to identify probabilities. No analyst can ever be “certain” of an outcome, but as I have said before, it is the technical analysts rather than the fundamental analysts who seem to have the least problems admitting to this fact. The market analyst in my opinion is like a bookie at the horse-racing track in that he is simply weighing up the various factors that will affect the betting odds for the horse. Previous form, ground conditions, weighting, etc., are all taken into account before the bookie posts his original betting odds and he will identify a favourite as the horse most likely to win. The technical market analyst weighs up price action, trend context, price derivative indicators, etc., before he settles on his favoured outcome to future market direction (his favourite horse!). However, like horse racing, the favourite does not always win and it is simply common sense to take this into account when analysing and, more importantly, investing in the markets.

Figure 5.4 shows the candle chart for the US Government Bond 10-Year Yield in the weekly time fractal. We can see that on 14 May 2004 there was a doji week which resulted in a seven-week trading range in the bond market. Previous to 14 May the market had been in a solid uptrend in yield (downtrend in bond prices) as can be seen from the long-bodied white candles, but then the doji week signalled that a lot of uncertainty had crept into the market psychology. The yield had reached a high of 4.90% but then market participants were encouraged to either stop putting large sell orders into the bond market or to cover their short positions, which meant that the yield stopped going up and resulted in the doji week. From there the market traded in a sideways fashion between 4.60% and 4.90% for a number of weeks – more than

patterns are doing is reflecting an analysis of the current market psychology and, therefore, they are no more than warnings or clues about a subsequent price action. So, when a doji appears it should, as always, be viewed in the context of the previous price action and it should not be assumed that it always means a reversal of any previous trend. Again this is where technical analysis methods can be seen interacting. A doji after a strong uptrend could just reflect the pause in the trend associated with a consolidation pattern in Elliott Wave or orthodox pattern recognition analyses. This happens many, many times as all the various tranches of technical market analysis interact with one another. Dojis often appear in the middle of strong trends and give rise to a period of consolidation before the trend resumes.

Figure 5.3 shows the Canadian dollar exchange rate to the US dollar in the monthly time fractal. The scale shows the amount of Canadian dollars it takes to buy a US dollar, and so we can see that the Canadian dollar (known as the “loonie” in the foreign exchange markets after the Loon bird that appears on the bank notes in Canada) was in a strong downtrend (uptrend in the chart) from 1991 through to 1994. Then, in April, May and June 1994 the market price action produced a series of three doji months. Now, in actual fact there was a small difference between the open and close in April (of 20 pips*) and May (of 10 pips) but, in the context of the monthly ranges between the high and the low (214 pips in April and 171 pips in May), the analyst can subjectively conclude that the price action is “doji like”. Again I realise that this subjectivity may not be to everyone’s taste but, to me, the candle analyst must always




Figure 5.3 USDCAD monthly – doji trading range

* A pip refers to the minimum price movement of foreign exchange rates. In this case the minimum price movement in the Canadian dollar exchange rate to the US dollar is 0.0001. The Japanese yen is quoted to the US dollar to two decimal places (e.g. 123.45) so the pip in that exchange rate would be 0.01.

is higher than the previous period's open. What is happening from a psychological perspective is that a previous downtrend is in place and the current period opens lower than the previous close, but this lower open is met with either capitulation selling or a surge in demand that causes a reversal in the supply demand mix. The market rallies throughout the period to end higher than the previous period's open, thus changing the bear's domination of the market. It is important to stress that this does not necessarily mean that the bulls will now be in total control of the market, only that the evidence suggests strongly that the bears have lost control of the market.

Now, although the level of uncertainty that the engulfing pattern is signalling in market psychology is not as great as the doji and, in the case of a bullish engulfing pattern, the odds on the bulls gaining traction are higher, it still represents an *exhaustion* of the previous downtrend and therefore acts as a brake in the first instance and not necessarily a quick shift into an uptrending regime.

The bearish engulfing pattern is exactly the opposite of the bullish engulfing pattern and has to occur after a preceding uptrend in the price.

Outside patterns are the westernised versions of the ancient Japanese price patterns and are slightly different in methodology. Rather than being the open and close levels that must engulf the previous levels as in the Japanese engulfing patterns, it is the highs and lows that must be "outside" the previous period's highs and lows in the westernised method. Although the pattern must occur after a preceding trend in the price, an interesting aspect to the outside pattern is that it doesn't necessarily matter if, in the case of a bullish outside pattern, the close is above the opening level. Yes it helps and it definitely adds to the confidence level of the pattern recognition if the close is above the open, but it doesn't have to be because the analysis of the change in psychology that is represented by an outside pattern is slightly different from that of a Japanese engulfing pattern.

Take the case of a bullish outside week pattern. The market has been in a downtrend for a number of months but on the week in question the market moves down and makes a new low below the previous week's low, but also during the week has moved up to make a new high above the previous week's high. The order in which the new highs or lows were made doesn't matter because this price pattern is not so much representing the bears losing control to the bulls as representing much more of a state of mass confusion. The market has made new extremes in both directions and so stop orders will have been taken out for both bulls and bears. That is, entering the week, a lot of the bulls would have gone long with a stop below the previous week's low and the bears would have gone short with a stop above the previous week's high. If the bulls get taken out they would be expecting the market to continue to go lower and if the bears get taken out then they would expect the market to continue to go higher. Alas, both bulls and bears get taken out together that week with no follow through, and this leaves both sets of participants agitated and confused. When market participants become more extremely confused than they normally are they become less confident, and when they become less confident they tend not to put too much risk on. When they do put risk on they tend to be quick to take it off again. This, of

course, is classic range-trading behaviour and it is why after an outside pattern the previous trend tends to exhaust and the market enters a period of range trading.

Figure 5.5 shows the Japanese candlestick chart of Honeywell International Inc., the diversified technology and manufacturing company in aerospace and power generation among others, using the weekly time fractal from January 2007 until April 2008. We can see from the chart that the share price of Honeywell was in an uptrending market – which it had been from mid-2006 – that accelerated rapidly in April 2007 when the price moved from \$46 to \$59 over the space of five weeks. The market then consolidated for a few weeks before, on the week of 22 June 2007, it gapped up to \$59.24 at the opening bell on Monday morning above the previous week's close. However, it could not sustain those gains and fell back. The market participants then took the price down to below the previous week's opening level and Honeywell eventually closed the week at \$55.68.




Figure 5.5 Honeywell International Inc. weekly candle – January 2007 to April 2008

This price action represents a weekly bearish engulfing pattern because the open was above the previous week's close and the close was below the previous week's open. That is, it was a down week when the previous week was an up week and the open and close levels engulfed the previous week's open and close levels. After an uptrend lasting several months, this price action screams at the trading regime analyst that there is a strong possibility that the uptrend is coming to an end and that either (a) the market will reverse its trend and/or (b) the market will enter into a period of range trading. What has happened is that the psychology of the market has changed. The market gapped higher on the Monday morning but failed to maintain those gains on account of either heavy supply coming into the market or demand tailing off rapidly. With the market closing below the previous week's open, the bulls would be seriously questioning the sustainability of a bullish view. The bears meanwhile

would have been unsettled by the fact that the market had gapped up at the opening on the Monday morning. So with both bulls and bears feeling deflated and therefore not feeling confident about trading Honeywell in any sizeable amounts, the odds on a change in the trend direction (either down or sideways) will have increased. Therefore at this point the analyst would be raising the probability of a range-trading regime developing.

Interestingly, the next week the market rallied a little and then advanced for a further three weeks before, on the week of 27 July 2007, the market experienced another bearish engulfing week. Another bearish engulfing week coming so closely after the previous one is further evidence that the underlying psychology of the market is no longer bullish. It might not be outright bearish or neutral but the confluence of bearish engulfing patterns suggests that the bulls in the market no longer have the upper hand. We can see that the market subsequently traded in a well-defined trading range for a number of months after these two bearish engulfing patterns appeared giving the analyst or trader many opportunities to take advantage of this slowdown in volatility (or range contraction).

Japanese candlestick bearish engulfing patterns and any candlestick analysis that involves an examination of the opening price in relation to the previous closing price is most valid when there is a gap between trading sessions or bars. If there was not a gap in time between the close and the open, or if the exchange rules mean that the price data will always mean that the open will appear to be the same as the close (as is the case with some cash indices), then the analyst has to be very aware that price patterns, like engulfing patterns, will be invalid. In the FX markets there is no gap in time between the close on one day and the open on the next day during the week because it is a 24-hour market. This is why I much prefer to look at weekly charts in the FX market when examining for engulfing and other such patterns that require a time gap between the close and open. Daily stock charts are fine because there is not a 24-hour market and the open can print well away from the previous close on certain occasions.

MORE CANDLESTICK UNCERTAINTY

The doji and the engulfing pattern are just two of many candlestick patterns that can signal trend exhaustion and uncertainty. Another three (of the many) that we can define here would be the hammer bottom, the shooting star and the hanging man. All of these candlestick patterns are strong indications of a change in psychology; of a trend exhausting and potentially turning to range trading.

Hammer Bottom

The hammer bottom occurs when, after a downtrend, the market opens near to the previous close, drops a lot during the period in question before closing the period up

towards the level at which it opened. The psychology here is that the initial selling activity that caused the market to sell off has abated and, indeed, been overwhelmed by relative buying activity in such a way as to cause the market to reverse back up before the close of the period in question. Coming after a downtrend it signals that the bears have lost control of the market. Crucially for our trading regime analysis, though, it does not automatically assume that because the bears have lost control the bulls have taken over. Again, I think that we must remember this when using candlestick analysis. Yes, a hammer bottom can mean that there is a sustained reversal in the market price but more often than not it is only signalling an end of the downtrend when the next period will be characterised by range trading.

Shooting Star

The shooting star can be considered to be the opposite to a hammer bottom and it occurs when, after an uptrend, the market opens near the previous close, rallies a lot during the period in question before closing the period down towards the level at which it opened. The psychology here is that the initial buying activity that caused the market to rally has abated and, indeed, been overwhelmed by relative selling activity in such a way as to cause the market to reverse back down before the close of the period in question. Coming after an uptrend, it signals that the bulls have lost control of the market but again it does not automatically assume that because the bulls have lost control the bears have taken over. More often than not it signals that supply and demand have become much more balanced than they have been of late and this balance can mean range trading.

Hanging Man

The hanging man is one of my favourite candlesticks, not just because of its evocative name, but because of the underlying market psychology it represents. The hanging man occurs after an uptrend in price when the market does not rise much but falls a lot during the period in question before finally closing back up near to the level at which it opened. The market psychology here is a bit more complex because the market crowd has driven the price down during the period but, due to the fact that selling has dried up and (or) buying has come in, the price has rallied back up before the close. You would think this is a bullish sign and at least a sign that the uptrend is still in force, but that is an illusion. As with all candlestick analysis the context of where they appear in the overall structure of price action is of vital importance. A hanging man has pretty much the same price action as the hammer bottom, but the psychology of the price action changes depending on whether the candle appears after a downtrend or an uptrend. Coming at the bottom of a downtrend it signals that the bears have lost control and, in fact if it appears as the market slowly rises

from that bottom it could be construed as a sign that the bulls are starting to take over. However, at the top of an extended uptrend, price action such as this when the market drops but rallies up again before the close can, on one level, be construed as the psychology of the “tail end Charlies” getting into a bull market *after* it has already done it. That is, the price action represents the last buyers getting into the uptrend and the last vestiges of ultra-bullish sentiment.

SUMMARY

In this chapter we have examined how an appreciation of Japanese candlestick analysis can help in identifying the transition from a trending regime (characterised by quite certain, strong psychology) to a range-trading regime (characterised by quite uncertain, weak psychology). Undoubtedly, Japanese candlestick analysis is invaluable in combination with other forms of trading regime analysis because it so graphically represents the changing dynamic of crowd psychology. This chapter has just skimmed the surface of the subject and for an in-depth study I would thoroughly recommend Steve Nison’s books on Japanese candlestick analysis.

In the next chapter we can see how we can add another layer to the probability of volatility analysis.

6

Volume Considerations

In the broad field – that is, modern technical market analysis – volume remains quite an underutilised indicator of potential information. In the early 20th century it was used a lot more, it seems, with Edwards and Magee focusing on it in their original work, but for some reason in recent decades the importance of volume analysis has been diminished, perhaps in favour of more quantitative methods that rely on price and price derivatives more than anything else. I think this is quite perplexing because there is a lot of informational content in volume analysis when combined with an analysis of price action, and it would seem to me that if anything was set up for high-speed, mechanised quantitative analysis then that was. My guess is that some of the most successful hedge funds in the world have some sort of high-frequency analysis of volume in relation to stock price movement in order to scalp the very short-term market, but they might be reluctant to divulge these methods for fear of this arbitrage being diluted. The more orthodox technical analysis of markets will look at volume on a discretionary basis, and this is what we shall concentrate on in this chapter in relation to how volume analysis can help us to determine the current and potential future trading regime.

Volume is the very essence of the market. If the market price represents the motor vehicle then volume represents the horsepower of the engine. The market cannot move in an extended fashion one way or the other without there being an imbalance of volume between buyers and sellers, but the market can stay absolutely still if there is massive volume, provided there is balance in that volume. When we analyse volume we are really getting into the nuts and bolts of the market's psychology. Essentially we are adding another layer of understanding of the market's psychology to just an analysis of market price. As an example, consider a situation where the market is at an important support level. If the market slips below that level the analyst might start to factor in that the market will continue to collapse below the support zone. However, what if the analyst noticed that the fall below the support level was being done on very small volume? Would he be quite as confident that the

support zone was really being broken? After all, if only a few people are selling, then maybe there is a chance that there is not as much horsepower behind the market move and that the price could be about to turn round and revert back to where it came from. On the other hand, if the fall below the support zone was done on larger volume, then maybe the analyst should conclude that there has been a lot of selling and that this might be the beginning of a trend.

INCREASING VOLUME ON BREAKOUTS

Figure 6.1 shows the OEX Stock Index in the United States of America, which constitutes the S&P Index of 100 highly capitalised stocks for which options are listed. This is the bar chart of the index using the monthly time fractal and the lower chart shows the monthly volume that was traded in the index using a histogram format. The scale on the volume chart is in billions of shares. We can see that the market had been trading sideways throughout 1993 and 1994 but then, in February 1995, it began to break above the sideways consolidation it had been in which, incidentally, happened




Figure 6.1 OEX monthly with volume

to be a symmetrical triangle. For the three months prior to February volume had been falling, but then in February, when the market price broke up past previous highs, the volume stabilised and was, in fact, a fraction higher than January's volume. This was the first sign that the breakup in price had the potential to be sustained. If the breakup in the price in February had been done on even lower volume than the previous few months, then it would have been a good sign that the breakup was to be treated with a healthy dose of scepticism.

The next month, March, saw an explosion in volume which was accompanied with the price rising even more and, at this point, the trading regime analyst would be very confident that the old range-trading regime has now been broken and that the market is unfolding into a trending regime. The increase in volume tells us that more shares have been traded, but when we combine this known fact with the other fact of what the price has done, then we can increase our understanding of the underlying psychology of the market. If more shares have been traded *and* the price has gone up then we can conclude that more demand is coming through for the market in question. This is a simple deduction given that the price will only go up if there is more being demanded than the available supply; therefore, if more shares have been traded than previously and the price has gone up, then it is reasonable to conclude that demand is on the increase.

In April we can see that the volume dropped, but in May it exploded again as the price rose once more giving more confirmation to the analyst that demand is rising as the price rises. With demand rising as the price rises having come out of a period of range trading, the odds are vastly increased that the market has entered a new trending period and this allows the investment manager to adjust his parameters in such a way as to take advantage of the trending regime. As can be seen from the chart, the OEX market did continue to trend strongly for quite a number of months after these observations.

VOLUME SPIKES AFTER TRENDING MOVES

As well as using volume analysis to signal the start of trend-trading regimes it can also be used to signal potential range bound markets. As has been mentioned in the section on behavioural analysis, the majority of market participants tend to buy near the top and sell near the bottom. This is just simply the way the market cycles in terms of psychology with fear and greed being the famous main drivers at such times. After an uptrend in which the majority of the market will not have participated by being long or overweight there will tend to be a rush into that market over a finite time scale and then, when all the buying has been done and there is no one left to buy, the market will fall. After a downtrend in which the majority of the market will not have been short or underweight there will tend to be a rush out of that market over a finite time scale and then, when all the selling has been done and there is no one left to sell, the market will rise.

These are both examples of capitulation by market participants who have been acting like giant pressure cookers sitting on their positions (underweight or overweight) and doing nothing while the pressure steadily builds. Finally when the pressure reaches the extreme they let off the steam by succumbing to capitulate and buy if they have been underweight or sell if they have been overweight. We have all been these soldiers I am sure, I know I have. It's the most difficult, gut wrenching, horrible feeling in the world when you are underweight and the market is roaring higher or you are overweight whilst the market is tanking. That is one reason why I now try to avoid those situations as much as possible by being a trend follower in principle, and so at least I can be safe in the knowledge that if the market makes a massive long-term move I shall not be one of the people experiencing those emotions.

It cannot be stressed too much that an analysis of volume must be used in conjunction with an awareness of what the price has been doing and that the volume analysis must be placed in the context of the current price action. When there is a volume spike *after* a trending period it is usually a good sign that there is a lot of capitulation going on. The post trend volume spike is the manifestation of the pressure cooker and when I see it in the charts I cannot help but think of that analogy and how relieved the people who have finally bought or sold must now be feeling! However, in the simple supply and demand world in which we live we know that if there is no one left to buy or sell then the market must correct the other way, at least slightly before it finds a new equilibrium.

A classic example of volume spikes signalling capitulation and either trend exhaustion or range trading comes when we analyse the western stock markets in the bear market run from 2000 to 2003. If we take the Standard & Poor's 500 Index of US stocks we see a situation where the capitulation psychology almost jumps out of the chart and slaps the analyst round the face.

Figure 6.2 shows the bar chart of the S&P 500 in the weekly time fractal with the associated volume histogram in the chart below. We can see that the market peaked in March 2000 and by December 2000 was in a well-established downtrend. Volume was beginning to increase during this period but it was a steady increase reflecting the fact that selling was increasing relative to demand. Increasing volume as the price is dropping is a good sign that the downtrend should be trusted. In March 2001 the market bottomed around 1080 on the index and volume at that point was the highest it had been. The week prior to, and the week of, the bottom in March 2001 volume spiked mildly and, from that analysis, the trading regime analyst should be on alert for a potential end to the trend, at least in the short term.

One of the most important aspects of market analysis and trading regime analysis in particular is to look for supporting evidence for conclusions from as many fronts as possible. In this vein, when we are looking at volume analysis, and volume spikes in particular, it can be very useful to combine this with an analysis of the Japanese candlestick chart patterns. We can see from the chart that the weekly candle for the week of 23 March 2001 was a strong hammer bottom with the market having a long spike down during the week but finishing the week back up near the open. This can




Figure 6.2 S&P 500 weekly and volume

be a sign of two things, both related. Either the market has gone down and found very strong demand from where the market has rallied, or the market has gone down on large, capitulation selling and, when all that has been done and there is none left to sell, the market has rallied on the supply demand imbalance. In many respects it is the supply demand imbalance that is the important aspect and not whether there was strong demand when the market spiked down. If there was strong demand *relative to the supply* then that is all that matters in order to push the price strongly back up, no matter if the nominal level of demand was not large. If the sellers have stopped selling then, even if there is not a big increase in demand, the price will stop falling.

So in March 2001 the S&P 500 index had a decent spike in volume and this, combined with the observation that there was a weekly hammer bottom, would have given the trading regime analyst a good indication that the downtrend was mature and that a period of range trading might well be on the cards or, at the least, the probabilities of a continuation of the downtrend were much lower than previously. We can see that the market did reverse over the next few weeks before continuing the downtrend in earnest.

Then in September 2001 the market had a huge spike in volume as the price collapsed down past 1000. Now, as this analysis of price and volume comes under the banner of technical analysis, we should really only be concerned with the price action and the related volume but it is worth pointing out (although I am sure most readers

will realise instantly) that this price and volume action was happening around the time of the 11 September 2001 terrorist attacks in the United States of America. We can see from the volume chart that the week of the attack resulted in tiny volume. A big reason for this was the actual physical reason that people could not get to work in New York and the exchanges were closed for a couple of days, but there was also a feeling within the global markets community that no business, certainly no speculative business, should be done as a mark of respect to those affected. Moreover, the whole of the civilised world was in shock at what had happened. I remember that day all too well, as I am sure most people do. I was working in the foreign exchange markets and, like all other markets, volume just dried up as everyone watched the horror unfolding. The next day something very poignant happened, however, that went some way towards restoring my faith in humanity. The foreign exchange markets have a reputation (undeservedly so in my view) of being the “wild west” of capitalism where the competition for business and trading profits is so tight and fierce that business ethics and morals are quite rare. However, on 12 September as the London markets opened again, the word was put out among the foreign exchange markets that only business that had to be done (for example, asset trade settlement) would be allowed and that no speculative business would be tolerated. This was because the shock was so great that no one wanted to try to make any money from the potential volatility an event like that could create. There were actually a few firms that decided to try to do some speculative trades but they were told by the market-making banks, who normally would be extremely happy to do any business, that not only were they not going to be given a price for that business but that the bank would not be doing business with them again. The customer’s “line was pulled” because of their perceived immoral behaviour. I thought this was fantastic and, in a way, a great ploy to foil the terrorists who no doubt thought that the global markets would go into extreme volatility because of their actions. The volatility in markets is only ever going to exist if business is being done. Look at the chart of the US Dollar Index over that week and you can see that the volatility was minimal. Some people might say that volatility was kept under control because the central banks flooded the market with liquidity and promises of liquidity, but I think that the fact that not many people actually wanted to do any business contributed more to the lack of volatility.

Returning to the S&P 500 index chart, though, we can see that in the week following 11 September volume returned with a vengeance. However, the market had a big fall on this massive spike in volume and, coming as it did *after* a period when the market had already tumbled significantly, the analyst would have been alerted to the fact that this volume spike could represent capitulation and signal the end of the move and/or potential range trading to follow. As it turned out the low that week marked the low for a number of months and over the next few months the market did generally fluctuate in a range. However, by May 2002 the market had begun to fall quite hard again and this trend lasted until July 2002 when, as we can see from the volume chart, there was another two-week period when volume spiked up to the levels seen in September 2001. In the second of those weeks, the week of the 19 July, the

market price action exhibited another classic hammer bottom in Japanese candlestick terms with the market opening around 847, falling to a low of 775 and then rallying all the way back up to close at 852. This was further evidence for the analyst that the volume spike resulted in a big imbalance between supply and demand and could therefore be assumed to be the result of capitulation selling. This price and volume action turned out to be the start of the reverse head and shoulder pattern that would eventually lead to the big turn upward in the US stock markets.

Volume spikes after a trending move has already taken place are a good indicator that the trending period is mature and that either the trend will reverse or the market will go into a period of range trading. In this regard volume spikes are a good measure of trend exhaustion and, if a trend is exhausting itself, then the market can do one of two things. Either it could reverse sharply into another trending phase or it could contract into a range-trading phase before continuing with the original trend or reversing it. More often than not, the market has to go through a period of replenishment (range trading) before deciding the direction of trend to take next.

As with other trading regime analysis methods, volume analysis can be applied to any time fractal because the psychology is the same no matter what time period we are looking at. If there was a big move in a stock on an intra-day basis and there was a big increase in volume at some point *after* the move had been going for a while, then the day trader could make an educated guess that exhaustion in the move is close. Similarly, an institutional investor whose investment time period is much, much longer should be keeping an eye on the monthly time fractal volume for clues as to potential trading regimes.

Figure 6.3 shows the bar chart of the US pharmaceutical company Merck & Co. (MRK) in the monthly time fractal with the associated volume histogram in the lower chart. Merck reached over \$91 in the mega bull market up to 2000 but it then started a precipitous fall until July 2002 when it bottomed at \$36.43 in what, at the time, was not only a volume spike month but also a classic hammer bottom month with the market moving down a long way during the month only to close the month back up towards the opening level. This price and volume action led to a period of range trading for at least a year when the market fluctuated between \$40 and \$60. Eventually the share price broke lower again in September 2004 and fell to a low of \$25.60 in November of that year.

We can see from the chart that the big fall in September produced an increase in volume but, importantly, the share price closed down very close to its lows on that month. So while this increase in volume on a large move down had the potential to be capitulation selling, the fact that the market had not recovered into the close should be telling the analyst that *perhaps* all the selling wasn't exhausted yet. So it proved in dramatic fashion with the next two months, producing off-the-scale volume spikes. In October the close for the month was off the bottom but not by much. However, in November, when the volume spiked even higher the closing price on the month was quite a bit off the low set during the month, making it a hammer bottom in Japanese candlestick terms. So at the end of November 2004 the trading regime analyst should

7

Previous Highs and Lows

Orthodox price pattern recognition analysis provides us with a methodology of using support and resistance areas or clusters in order to identify patterns in investor psychology and to project into the future how these patterns might be resolved. A lot of these patterns use horizontal support and resistance levels to identify the zones where the boundaries of investor psychology defined by price action may lie. However, horizontal support and resistance zones can provide further assistance in analysing potential trading regimes.

Previous high and low points act as terms of reference for everyone involved in the markets. The financial press are constantly coming out with headlines such as “Gold Makes New Historic High” or “The Dollar Makes New Historic Low”. As has been said before, it is the prices of markets that make the news not the other way round! However, more than this, investors and market participants are mostly aware of previous highs and lows in stocks, bonds, currencies and commodities on either a real or nominal basis. I suppose it is one of the basic valuation techniques, especially when it is used on a real (adjusted for inflation) basis. As I mention later, I do not agree with those who talk about market behaviour in terms of people being aware of certain price levels as being what they would call a “self-fulfilling prophecy” because the changes that occur in the price trend can persist for a long time and, if only a minority of market participants were focusing on a particular level, this persistence could not occur. Nevertheless, it is this awareness of previous highs and lows that can directly affect, at least some, investors’ behaviour when those levels are approached.

Let’s say a stock had a very good run a few months ago and the share price reached a level of, for the sake of argument, 50 before falling back, over the course of a number of weeks, to 25. Then the market picked up again and the share price began another bull run going through 30 and 40 before coming up to the 50 mark again. What do you think the psychology of investors in this market could be? Well that depends on a multitude of factors but let’s assume for simplicity that whoever was long this stock was long in quite a number of shares. This is just basic common




Figure 6.3 Merck & Co. monthly and volume

have been concluding that Merck had experienced a massive amount of panic selling and that, very probably, the selling had been all but done. Therefore the probabilities vastly increased that the downside risk to the Merck share price had diminished a great deal and that the market was now either going to snap back up or more probably enter a period of general consolidation and range trading. So it proved with the market fluctuating between \$25 and \$35 over the next 18 months. Eighteen months of range trading is a long time even by the time scale of institutional investors and so there were many opportunities to factor this into an investment process and benefit from the analysis by either trading the range via some oscillator techniques, by selling volatility via options or by a carry strategy of collecting the dividend income in a stable share price environment.

SUMMARY

This chapter has shown how factoring volume analysis into our process can help us to determine the probability of the next trading regime. The next chapter examines a simple, yet powerful observation that can also add to our armoury.

sense because the market could not have rallied from the lows at 25 without a lot of share buying. So, our stock is approaching the 50 level again, which was the previous high point. Is it fair to assume that the person or people who had bought the stock on the run up from 25 and had held it would now be thinking “Well, I have had a good run and we’re now coming up to 50 again. That was the previous high point. Maybe I should sell out at least some of my holding.” In my opinion it is very fair to assume that because human psychology is such that as soon as profits are shown in an investment then the old devil appears on your shoulder whispering into your ear “Close the position. Don’t let your profits turn into a loss.” This pressure builds and builds the more profits are shown until finally people succumb and close at least some of the position. So if, in our example, people do succumb to the temptation of selling some out because it is near the old highs, this influx of supply into the market will no doubt create, at least in the short term, an imbalance between supply and demand. The market should at the very least stall around those old high levels due to this investor behaviour.

Another group of people to think about in this example is the people who bought the shares near the top of the previous cycle. That is, they bought it the last time it got to 50. In a way these people are more important to an understanding of psychology around previous highs and lows than the group mentioned previously because their emotions might be even more powerful than the others. You know the feeling. We have all done it. We buy it near the highs and watch it come off almost immediately. We don’t sell because we would take a loss and we think it will come back. It doesn’t. Instead it languishes, well below our buying level for a long time before finally turning back up again. Rather than want to buy more at this point, we have an overwhelming feeling of relief as the market comes back up to our buying level and we think to ourselves “well, if I can get out close to my in price then I will be more than happy”. So we sell out when the market returns near to our in price which was the previous high and that psychology behind supply, if it is replicated, can cause an imbalance between supply and demand dynamics that result in an exhaustion in the price trend. This psychology is particularly relevant when the previous high (when we and many others would have bought) was very brief in terms of time. That is, the market went quickly up to 50 the first time but didn’t stay there for very long and therefore created a situation where people who had bought on the spike up were trapped and couldn’t get out without taking a big loss.

It is with these psychological factors at work on supply and demand in the market that previous highs and lows can and do act as zones of congestion, where the price can range trade or at least exhaust after the previous trending move.

Figure 7.1 shows the bar chart of the exchange between the Norwegian krone and the Swedish kronor in the weekly time fractal. This exchange rate trades with quite low volatility normally and a lot of that is due to the fact that the actual or perceived differences between the two economies at a macro level tend to be quite small. However, the rate can go through strong trending periods and one of them is shown here in the period from mid-2000 when the Norwegian krone started to




Figure 7.1 NOKSEK weekly – old high range

rally very strongly versus the Swedish kronor, moving from around 1.00 to a high of around 1.26 in September 2001. The market then exhausted and fell back down to a low around 1.1470 in January 2002 before starting another rally back up. It took a few months, but the market eventually traded up to and just over the previous high of 1.26 in July 2002. At that point the market failed to break decisively through into new high ground with the candlestick chart on the week of the 12 July showing a shooting star formation, where the market goes quite a lot higher during the week but fails to hold the gains and ends up back down towards the lower end of the range.

This is a sign of a good supply/demand imbalance on top of the market, perhaps because all the buyers have bought or that lots of selling has come into the market at that level. Given the fact that the price pattern was taking place around the previous highs it is a fair guess to assume that people who had been long the market were taking this opportunity of being back at the old highs to sell out at least some of their Norwegian krone. It could also very well be that some longer term investors who had bought “Noky Stocky” (Nok for Norwegian krone and Stocky, Stockholm, for Sweden) around the previous highs back in September 2001, but had been holding on to the position as they watched it fall for the few months thereafter, were now feeling very relieved that the market had managed to regain their buying levels and that this would allow them to get out of the position fairly flat. Notice that the previous high in September 2001 was quite spiky in that the market went quickly up to 1.26 but didn’t stay there for too long at all before coming quickly back down. This is the type of price action that creates trapped bulls whose only thought afterwards is “what is the best price at which I can get out of the position?”

The previous high in this example was acting as a step-off point for people and the trading regime analyst, witnessing this, should be alert to the possibility that this

cheer went up from that part of the floor. This highlighted the fact that the focus for the market at that time was 10,000 and for the months before that the market “chatter” was about reaching 10,000 on the Dow. We can replace “market chatter” with market psychology because the market chatter just reflects what people are thinking and if what people are thinking translates into what they actually do in the market, then the expectations were for 10,000 to be reached on the Dow. People who were already long could not add to the demand, but people who were not would have been affected by this psychology and possibly bought. The fact that the market continued to rise is evidence that this was the case. So when the market finally touches 10,000, having had a strong multi-year rally up to that point, should we be surprised that the pace of the rally begins to slow down dramatically and, in fact, the market then finds it difficult to pull away from this zone? A lot of supply will have come on the market as market participants who were long the Dow from a few years ago would have thought (subconsciously maybe, but they would have been aware of it) that 10,000 on the index felt like a good place to book some profits on those long positions. This is simply human nature and no matter what people tell you about the fact that it wouldn’t enter into their investment process because they look at things objectively, such as P/E ratios and earnings data, etc., at the back of their mind is Dow 10,000 nibbling away at their thought process. Does this affect their actions in the market? Of course it does.

A few years ago a very well respected fellow technical analyst, John Breame, pointed out to me an interesting observation, that when the Dow Jones Industrials hit 100 for the first time in 1906 and it took another 19 years to get significantly away from that zone. In other words, there was a big trading range with 100 (or thereabout) acting as the upper limit over a multi-year timeframe. We can see this in Figure 7.3 and, of course, once 100 did break, the resulting rally in the Roaring 20s was very powerful, and almost as powerful as the crash that followed it. However, the point is that 100 acted as a general resistance (or attractor) zone for a long time.

In 1966 the Dow Jones Industrials hit 1,000 for the first time and it took another 19 years before the market moved significantly up past that zone. Again the market traded in a big trading range with 1,000 acting as the upper resistance zone. Yet again the market attractor seemed to be a big round number with psychology no doubt affected when the market got there.

At the time of writing, March 2008, we have already seen the Dow in a big trading range since it hit 10,000 in 1999. The market dropped to around 7,500 but did rally above 10,000 to 14,000 before dropping back again, so admittedly it is a wide trading range but essentially the Dow has been very flat in nominal terms since 1999 (in real terms it has actually been in a multi-year bear market, but most market participants focus on the nominal Dow). With nine years down, if history is to repeat itself, then we might be in for another 10 years of big range trading before the Dow finally clears above the 10,000 zone in any meaningful fashion. I suppose the good news might be that if the logarithmic sequence is to repeat itself then the next secular bull market should take us roaring up to 100,000 on the Dow!

a slightly higher price. As MCD approached the previous high during August 2005 the analyst should have been assuming that there was a good potential that a group of market participants had an overriding psychology of getting out of the stock with as little pain as possible.

So in the week of 12 August as the share price of MCD went just over the previous high before falling back into the close, the trading regime analyst should be giving quite a high probability to the chances that the up move in MCD was exhausting, and that either a fall or a period of range trading was about to occur. In the event, we can see that from that high during the week of 12 August the market did indeed trade in a very distinct sideways movement for a number of months thereafter (in fact for nearly a year), allowing the trader or investor to factor that decline in volatility into his trading or investment process.

SUMMARY

In this section we have seen how zones of longer term significant previous highs or lows can act as areas where the psychology of market participants can change from being confident to less confident about market direction and therefore from the market being in a state of disequilibrium (or trend) to a state of equilibrium (or range). In preceding chapters we have examined the patterns of price and volume that appear as a direct result of the psychology of individuals acting within a crowd and how they are the basis of the volatility cycle we witness in the markets.

In the next section we will look at a phenomenon that, although *extremely* simplistic, is nonetheless a recognised behavioural concept in the markets that can help us in our analysis of the potential market trading regime.

THE LURE OF THE BIG ROUND NUMBER

One of the most interesting aspects of market behaviour is that round numbers seem to act as attractors for market participants. What I mean by this is that quite often the market will stop a move when it reaches a round number (such as 10 or 100 dollars on a stock price) and it will either reverse sharply or enter a period of range trading. When we think about it, we should not be surprised that this is the case because the markets, as mentioned previously, are merely a collection of human beings buying and selling stuff to each other and are therefore driven predominantly by the psychology of each human being aggregated up to the crowd level. Round numbers ignite our imagination and are easier to think about. Ten thousand on the Dow Jones Industrials index will get more prominent headlines in the newspapers and other media than 9,563. I remember vividly when the Dow hit 10,000 for the first time back in March 1999. I was working on the same trading floor as equity investment managers and dealers, and when the index finally printed 10,000 on 16 March 1999 a huge

overhead supply in the market may well persist for some time leading to potential range trading for the cross rate. This turned out to be the case here as the market then traded in a very distinct horizontal fashion for a number of months afterwards, allowing the trader or investor to take advantage of the declining volatility.

Figure 7.2 shows the bar chart of the share price of McDonald's Corporation (MCD) in the weekly time fractal from 2003 to 2006. We can see that the market bottomed in March 2003 around \$12 and began a strong rally over the next couple of years, culminating in a high of \$34.56 in March 2005. The move up to the high in the previous few weeks was relatively fast and by the end of the week of the high it was back down to \$32.50. The share price continued to fall over the next few weeks down to a low of around \$27 in July 2005. Then the market started another rally until, in the week of 12 August 2005 it reached the previous high level of \$34.56. That week it reached a high of \$34.70 before falling back into the close at \$33.26.




Figure 7.2 McDonald's Corporation weekly – old high range

Now, this price action wasn't a classic shooting star in Japanese candlesticks but the price action psychology of the market going up just past a previous old high and closing the week a reasonable amount below it again should not be considered to be particularly bullish. It might not be bearish, but it should not be considered to be bullish. So this was one piece of evidence that the trading regime analyst received from the market regarding the potential future trading or volatility regime. Another bit of information was that the previous high in March 2005 was quite short lived and so the analyst can assume that there might have been a lot of people who had a rush to buy MCD in the quick run up to that high only to find that they were the last buyers. An analysis of human nature would tell us that a large percentage of these buyers would not have sold MCD as it dropped hard in April through to July, but will be waiting for the opportunity to sell as close as possible to the price they got in, or at




Figure 7.3 Dow Jones Industrials yearly – logarithmic scale

The observation of the Dow pausing and range trading at big round numbers like this opens up an interesting and related issue. As mentioned earlier, Robert Prechter from Elliott Wave International has, with the help of others, created a movement known as Socionomics. This is an incredibly interesting new development in market analysis because it would seem to be the link between what we might commonly think of as fundamental analysis and what we think of as technical analysis. For me the most interesting aspect is the fact that the socionomic way of thinking turns orthodox thinking on its head. For instance, one well-quoted example of the way socionomic thinking differs from orthodox thinking is that an orthodox economic thinker would probably say that a recession makes businessmen cautious, whereas a socionomic thinker would say that cautious businessmen make recessions. Another one might be: Does war make people angry or do angry people make wars? It is this cause and effect aspect and how it relates to human behaviour and psychology that I believe is truly fascinating, and socionomics is incorporating these issues into traditional economic analysis. What socionomics is doing is studying the dynamics of social mood (as measured by the stock market which is the barometer of social mood) and relating that to developments and attitudes in society. What drives what? Does social mood drive the economy or does the economy drive social mood? Is society more liberal and tolerant when the economy is doing well and less liberal and tolerant when the economy is doing badly? The evidence for this is there for all to see, and I am not going to replicate it here because Elliott Wave International and the

Socionomics Institute do such a great job with it. More information on socionomics can be found at www.socionomics.net.

As the stock market goes so goes the economy, and when the Dow hits the big numbers like 100, 1,000 and 10,000 in the examples above, the socionomic aspect we should be thinking about is whether it is (a) the subsequent economic stagnation that is making the stock market stop there for such a long time or (b) people becoming cautious at these numbers in the stock market that contributes to a period of stagnation in the economy?

Figure 7.4 shows another example of round numbers acting as attractors and potential range-trading environments, this time in the Canadian dollar. The Canadian dollar (or the “Loonie” as it known in the FX markets on account of the bird, the Loon, on its banknotes), has for years been derided by the Americans as the “northern peso” in reference to the fact that it was perennially weaker than its close neighbour, the mighty greenback American dollar. Then the Canadians opened up the Alberta oil sands. Well, that was just one particularly bullish bit of news for the Canadian dollar amid a plethora of both bullish and bearish “news” that has transpired over the last few years, but whatever the “news” and reasons for the Canadian dollar rallying, the price history tells us that the exchange rate of the Canadian dollar to the US dollar did fall quite rapidly from a high of 1.62 in 2002 and accelerated down from early 2007.




Figure 7.4 USDCAD weekly

The last time the USDCAD exchange rate had been at parity (i.e. 1.00) was in the early to mid-1970s and so the mentality of the market participants as the market fell from 1.20 was one of growing excitement that parity (the symbolic parity level) would be seen. I remember that when the exchange rate had fallen only to 1.10,

one bank in particular had started to organise a “parity party” for its clients! The psychology was incredibly bullish for the Canadian dollar and of course sometimes that can result in a big reversal, depending on how many people not only hold that view but also crucially hold that position. After a long pause that started with this move to 1.10, the market moved very quickly lower, blasting through the parity level for a few weeks before dramatically reversing in November 2007 from a low around 0.91. The market then moved higher to the parity level by the end of November and then traded around that level in a distinct sideways range for a long time. At the time of writing the market was still trading in this range. We can say, therefore, that, notwithstanding the “overshoot” when the parity level was first broken in 2007, the parity level has acted as an attractor in the Canadian dollar exchange rate with the US dollar with a distinct range-trading regime taking place around it. Identifying this might allow the analyst or investment manager to factor this into his overall trading game plan either in terms of altering position size or choosing a trading technique that allows him to take advantage of range trading.

SUMMARY

In this section we have highlighted a simple and easily dismissed phenomenon of market behaviour that can add to the chances of trend exhaustion and range-trading regime developing. It may sound too simplistic to some and no doubt some people might point out that, to their minds, it is a self-fulfilling phenomenon but the lure of the big, round number is an all too present reality of the financial markets and therefore one that cannot be ignored.

In the next chapter we will find a theory, or principle, that is in many ways the glue that links these patterns in price and volume together and which could be considered to be the purest form of technical market analysis.

8

Elliott Wave Principle

*“For want of a nail, the shoe was lost;
For want of a shoe, the horse was lost;
For want of a horse, the rider was lost;
For want of a rider, the battle was lost;
For want of a battle, the kingdom was lost!”*

Ancient nursery rhyme

We now come to a method of market analysis that could be regarded as the closest one might get to being a complete investment method, that takes account of every time frame and every regime of market price action. It could easily be thought of as a complete trading regime analysis methodology because the philosophy behind it is rooted in the very same philosophy that is behind the concept of volatility cycles.

In the 1930s Ralph Nelson Elliott discovered that the US stock market displayed certain characteristics of repeatable and consistent patterns. He identified these movements as “waves” and noticed that these waves occurred at every time scale to form part of the next larger wave. What Elliott had discovered is something that certain Chaos Theory advocates, and statisticians would decades later try to claim was a new discovery – namely the fractal nature of financial markets. The contribution of Chaos Theory to understanding complex systems has, however, gone a long way to validate the discovery that Elliott made and to vindicate those who use and understand the theoretical underpinnings of Elliott Wave. What is interesting is that Chaos Theory backed up the hypothesis that complex systems undergo “punctuated growth”, meaning that systems will move and then consolidate, move and then consolidate. So what they are doing is forever building “fractally” into similar patterns of increasing size.

The Elliott Wave Principle could well be considered to be the definitive all-round trading regime analysis methodology. In fact this methodology, like orthodox pattern recognition, does attempt to anticipate the price *direction* but for the purposes of this

book we are only concerned with the trading regime or volatility analysis aspect of it. This methodology has more than its fair share of detractors and critics, but it also has just as many people who consider it to be the most useful, pertinent and theoretically sound trading or investment methodology that exists. One thing is for sure that in a discussion about the Elliott Wave Principle, the debate will become quite heated, not only among those people who really do not know what they are talking about when it comes to the Wave Principle, because they have not studied it in any depth, but also between those people who are familiar with the methodology but who argue on the finer detail of the analysis.

Technical analysis is the study of past price action in order to anticipate (not predict!) future price movements and it is rooted in human psychology and behavioural analysis. There are many, many strands of technical analysis but two broad methods within the subject are Elliott Wave Principle and orthodox pattern recognition. In common with the theme of technical analysis, both of these methods take the price to be the manifestation of all that is known and expected about the future at any moment in time. What sets these methods apart from a lot of technical analysis is that, at their purest level, they only consider the *price* of the market and are not concerned with derivative indicators like moving averages or oscillators. For these methods the price *IS* the sentiment indicator – if the price is going up, the market is bullish, and if the price is going down, the market is bearish. As Robert Prechter from Elliott Wave International puts it, the pure price of the stock market is the “barometer” of sentiment about the economy.

Now, a full detailed explanation of Elliott Wave analysis is well beyond the scope of this book so I will only give a brief overview of the salient points in terms of how they relate to trading regime analysis. For anyone wanting to learn more about Elliott Wave I would recommend Robert Prechter’s Elliott Wave International at www.elliottwave.com.

The basic Elliott Wave pattern consists of impulsive waves and corrective waves. There are five sub-waves in an impulsive wave and the impulsive wave moves in the direction of the trend of the next larger fractal size. A corrective wave has three sub-waves and moves against the trend of the next larger fractal size. Numbers denotes impulsive waves and letters denotes corrective waves. Figure 8.1 makes this fractal relationship clear and shows a classic stylised Elliott Wave cycle.

We can see that after five waves up (an impulsive move) the Elliott Wave analysis suggests a period of consolidation (a corrective move) in an A–B–C correction. This is, in effect, a stylised model of how price would behave in a volatility cycle where the impulsive waves would correspond to periods of increased volatility (range expansion) and the corrective waves would correspond to periods of decreasing volatility (range contraction). Note the word stylised! This is not the real world and one of the problems most people find with an Elliott Wave is how difficult it is to apply to real world price action. Real world price action generally does not look like this diagram, although sometimes it can come very close indeed. No, the reality we




Figure 8.1 Stylised Elliott Wave cycle

are faced with as market analysts is that we have to use the rules and guidelines of the Elliott Wave Principle to identify the best (or preferred) wave count at any point of time.

The three rules of an Elliott Wave are:

1. Wave 2 cannot retrace more than 100% of wave 1.
2. Of waves 1, 3 and 5, wave 3 can never be the shortest.
3. Wave 4 can never end in the price territory of wave 1.

And the three most important guidelines of an Elliott Wave are:

1. When a five-wave (impulsive) pattern is completed, the subsequent three-wave correction usually terminates near the extreme of the previous fourth wave of lesser degree.
2. If wave 2 of a five-wave (impulse) pattern corrects sharply against the trend, expect wave 4 to be a sideways correction.
3. If wave 3 of a five-wave (impulse) pattern is the longest, waves 1 and 5 tend to be equal in length.

If the analyst follows these rules and guidelines, then price action analysis using the Elliott Wave Principle becomes a fascinating and exciting occupation because the “anticipatory” nature of the work lets the price action unfold in front of your eyes. Moreover, these rules and guidelines keep the analysis honest and can stop the analyst from going down the wrong path with a count. However, even though these conditions would seem to permit a mechanical model of an Elliott Wave to be generated via computer, the reality is that due to the chaotic, dynamic and evolutionary aspect of the market price action the analysis is still, to this day, best done via a human analyst. The best in the business, Robert Prechter and Elliott Wave International, have been researching whether it can be modelled and programmed but they have still not managed to develop a model with which they are satisfied. If the recognised world experts on Elliott say it can’t be done, then that is good enough for me, even though that won’t satisfy the quants.

When one studies technical analysis, one eventually realises that there is a lot of overlap between the various strands. In particular, orthodox pattern recognition and the Elliott Wave Principle are very closely related. A fourth wave symmetrical triangle is an obvious example but flags and pennants are merely pauses (or sub-degree retracements) in a bigger trend. Outside days and candlestick reversal patterns occur quite frequently at Fibonacci retracement and extension levels of the previous price move. So much of the Wave Principle is found in other forms of analysis, even from a fundamental perspective. The personalities of the individual waves correspond so well to what is happening in the real economic world. Wave 1 of an uptrend takes place when all the fundamentals still appear bad. Wave 3 of an uptrend takes place when all the fundamentals start to improve. Wave 5 of an uptrend takes place when the improvement of the fundamentals are tapering off. Hence you have the inevitability of the smart money (the minority) getting in and out before the big trend ends, while the dumb money (the majority) get in on wave 5 and generally end up losing. This is human psychology at work. No one wants to buy when the news is all still bad (wave 1) but everyone wants to buy when the news is all good (wave 5). Why? Because human psychology dictates that we generally need reassurance and feel safer in crowds.

This is what I like about the Wave Principle. It has all the attributes of a complete trading methodology. As it is both trend following *and* contrarian, it can give us, as market analysts, a fair idea of the trading regime (trending or ranging) we might be in over the next period. It also has the other attribute that a lot of technical analysis does not have: it can be “anticipatory”. One of the frustrations of being a trend follower is that you don’t know when the trend will end and are apt to ride the tiger until it turns round and bites you. That is, you ride the trend using your moving average system or such like, but you know that you are always going to give some of those profits back when the trend eventually ends. With the Wave Principle, it is possible to get a decent idea of the level where the trend *might* end by using the rules of the Wave Principle and Fibonacci calculations. This does not mean that the trend always ends in such a calculated place but, by being aware of these levels, the market trader can use this in his overall money management plan.

On one level we could think of the Elliott Wave as a complete trading regime analysis methodology because of the fact that the impulsive waves should have a trend-like nature and the corrective waves should have a range-trading quality. Well, yes and no. Although the adjectives *impulsive* and *corrective* conjure up images of trends and ranges, in reality the “impulsive” waves do not all exhibit trend-like qualities and the “corrective” waves do not all exhibit range-trading qualities. However, and this is the most important point, generally they do exhibit these qualities and what we are looking for as trading regime analysts is probabilities. The one aspect about Elliott Wave Principle that many people dislike is its perceived subjectivity. Many times they will see an analysis that had one particular wave count only for the analyst to change the count to something different after the fact. This *alternate* wave count aspect of Elliott, where there is always a preferred and alternative count depending on how the price action turns out, can be so very frustrating and confusing

that many people give up on the analysis. However, in any form of market analysis, there are only ever probabilities and never certainties, so if we can add together pieces of evidence that suggest a trending regime is about to occur, or that a range-trading regime is about to occur, then we can anticipate, alter our risk parameters and see how the scene changes with this flexible game plan.

HAVING A FLEXIBLE GAME PLAN

Tiger Woods always goes out onto the golf course in a tournament with a game plan. In 2005 he played the British Open Championship at Royal Liverpool Golf Club when the course was so dry and burnt that the fairways had no green on them whatsoever, and this made the course play very “fast” with a lot of run on the ball after it landed. As a consequence he decided that the best game plan would be to try to keep the ball in play as a priority and not let any loose tee shots wander into unplayable places. He therefore took an iron (played a shorter more defensive club than usual) off just about every tee that week and his final round was described by Jack Nicklaus as one of the finest, most controlled rounds of golf he has seen. Woods won.

What if, after the first round, there had been huge amounts of rain that soaked the course and made it play differently from the previous conditions? Would Tiger Woods have changed his game plan? He may not have altered it radically but he would have tweaked it a *bit* to reflect the new conditions. This is what top sports people and top sporting teams do. They always, *always* go out with a preferred game plan and stick to it as much as possible but, if it becomes obvious that a new game plan is required, they do not hesitate to switch to “plan B”.

Having a flexible game plan is imperative in any competitive environment because you, the competitor, cannot control how the opposition or the conditions will turn out and, in my opinion, trading or investing in the markets is very similar to a competitive environment, especially a sport like golf. For those readers who are not golfers, I beg your pardon to bear with this section because I think it makes a lot of sense even to those who are not familiar with the game.

There is an old saying that golf is a game that can be played but never won. In other words, you can play the game of golf and enjoy it, but you will never ever conquer the game because there will always be days, no matter how good you are, when you will play poorly. Even Tiger Woods and Jack Nicklaus have had days and competitions when things just don’t go their way, and this, in the final analysis, is the inherent beauty of golf. Those of us who have played the game know that no matter how badly you play there is always that one shot that was hit well and you watched the ball fly towards the target with a great, great feeling inside of you. That is the shot that makes you come back to the golf course. It is the same in trading or investing in the markets. No matter how much of a poor run of form you are having there is always that one trade idea or, if you are systematic, model enhancement that has the potential for better returns. You know it has worked in the past so it should work in the future. Just like golf, the markets can “hook” someone and ignite a passion.

five-wave movement. Figure 8.2 shows the move higher and how the rules and guidelines of the Elliott Wave are met, resulting in the correct labelling of this move as a five-wave impulsive move up. So what does the trading regime analyst do now? Remember, as a trading regime analyst our goal first and foremost is to identify our current trading regime and when the probabilities are skewed towards a change in regime. By the end of February 2007, therefore, the trading regime analyst using Elliott Wave as a tool will have realised that the EURGBP market is in a trending, impulsive phase on the daily fractal but it has just completed a five-wave movement. This is valuable information because the analyst can then conclude that it is probable that the next phase in market price action will be a period of consolidation and a corrective move. Figure 8.3 shows that once the five-wave impulsive up sequence ended, the market indeed went through a period of range trading that allowed the investment manager many opportunities to take advantage of that range either by switching investment techniques or by reducing position size, or by selling volatility in general.

Figures 8.4 and 8.5 show another example using a longer time frame (or time fractal) in the Australian dollar exchange rate versus the US dollar (AUDUSD). In April 2001 the “Aussie” bottomed and began a long advance over the next three years. Whether the analyst was expecting this advance or not is not the question for us as trading regime analysts. Remember, what we are interested in is identifying the current trading regime and when it is likely to change, and for the purposes of using Elliott Wave analysis to support this aim we are really only interested in when we have seen a clear five-wave move. Thus, the impulsive nature of the moves higher in the Aussie in 2002 and 2003 would have made the analyst realise, if he hadn’t already, that the market was in a trending, impulsive trading regime. At this stage




Figure 8.4 AUDUSD weekly

Another similarity between golf and trading the markets is having the ability to accept a mistake and try to minimise the damage that the mistake may make. When you hit your drive into the deep rough, most of the time the best thing to do is to “chop it out” back on to the fairway where you have a good chance of just walking away from the hole with a slightly worse score than par rather than try to go for the hard shot to the green from the deep rough, when you could easily end up making a disastrously high number of shots. As the golfing expression goes, if you hit it into the deep rough you should “take your medicine and move on”. It is the same in the markets. Taking your medicine is analogous to taking a stop-loss on a trade, but in order to do that you must have the discipline and *humility* to admit that you had made a mistake in the first place. If you do not, and you continue to run the losing trade, then, as we have seen in the past, it can end in disaster.

Ah yes. Humility. If ever there was a personality trait that was required in golf and, more importantly, in the markets it is this one. Humility, the quality of being humble, is so vitally important because it is the emotion that allows the ego to be kept in check, and having an overactive ego is one of the most destructive forces in golf and in the markets. I remember reading an article about a CEO of a top US company who had a policy of not recruiting managers who did not play golf, not because of the popular misconception that the golf course is where business is done but because he believed that “people who do not play golf do not know what it means to be humble”. He wanted people in his company who would understand that the ability to admit mistakes and show that you did not have all the answers was a sign of strength and not a sign of weakness. I also remember having a conversation with the Chief Investment Officer from a company I had previous dealings with whose decision to remain overweight stocks in a deep bear market contributed greatly to an almost disastrous situation for the company. He told me on one occasion that he did not play golf, and I thought that made sense.

Now, we could spend an entire book discussing the Elliott Wave Principle but what I want to consider is how to use the Elliott Wave when analysing the potential trading regime in which we are likely to find ourselves. In this respect we can disregard a lot of the intricacies of the Elliott Wave, valuable though they may be, and concentrate on when the wave count gives us a high probability of a trading regime change. When I was working at Hambros Bank in the 1990s I worked closely with Tony Plummer, a recognised expert in wave analysis. Although his own theories on wave analysis differ slightly from Elliott’s, I will always remember something he said to me. We were discussing the intricacies of wave analysis and how there was always a preferred and an alternative wave count, when he pointed out to me that, whatever the difficulty in identifying what will turn out to be the correct wave count in the end, when you see five clear impulsive waves up or down then you know that the next phase will be a consolidation in the price action. I thought this was a very useful piece of information but it wasn’t until a few years later, when I started thinking more about trading regime analysis, that I realised just how valuable this insight could be. What the Elliott Wave analysis gives us is, in effect, a high probability trading regime

indicator for when the market goes from a trending (impulsive) regime to a ranging (corrective) regime.

When a price chart is showing five distinct waves higher, then it makes sense from an Elliott Wave perspective to expect a period of consolidation. This is consistent with the notion that markets go through periods of trending and non-trending (ranging) as the volatility cycle goes up and down in a mean-reverting fashion. Examples of this are shown in Figures 8.2 and 8.3. In January 2007 the EURGBP exchange rate started a move higher that, by the middle of March, was a very clear




Figure 8.2 EURGBP daily



Figure 8.3 EURGBP daily



Figure 8.5 AUDUSD weekly

the trading regime analyst who is using the Elliott Wave as a tool should be asking himself the questions of when this impulsive phase will end and the next phase (the ranging or corrective phase) will begin.

Counting Elliott Waves is not an exact science (something that the quantitative community have found very frustrating over the years!) because with markets we are dealing with a complex system that is under constant evolution and change. All we can do is deal in probabilities and, if the anticipated outcome does not come about, we use that price movement information to update our probable expectations. So when, in February 2004, we label the move up from 2001 as a five-wave sequence we don't obviously know for sure that the next phase will be a ranging or consolidation phase. But we have a good idea based on our understanding of Elliott Wave analysis that a ranging regime is, at *that* moment, more probable than not. As it turns out, in Figure 8.5 we can see that from February 2004 the Aussie dollar goes through an extended period of consolidation and we have a range-trading regime for about three years. The use of a trend-following strategy in that environment, depending on your parameters of course, would have underperformed a contrarian strategy because the trading regime has switched and the volatility cycle has turned.

THIRD OF A THIRD AND THIRD WAVES IN GENERAL

As well as using a completed five-wave pattern as a high probability indicator for a ranging regime, Elliott Wave analysis also gives us another high probability situation, this time for trending regimes. As we can see from the market psychology cycle in

Chapter 3, the turning point of markets, depending on what fractal we are concerned with, happens actually quite slowly. If the market has been in a downtrend, then the psychology of market participants will have to go through the despair (when the market bottoms), denial/comfort and happiness (when the market moves up a little but the popular psychology is still one of impending doom) and happiness (when it is seen that the market has not made a new low and, in fact it is continuing to rally). This psychology cycle has to take place before the market can *really* take off again in an upward trajectory and these stages in psychology lead to the price action of the markets.

If we take the example of the stock market, the despair stage is when most people will sell out having no doubt held on, hoping that the market would rally. This is the bottom. The large selling at this stage creates an imbalance between supply and demand in the market and so the next stage, the denial/comfort stage, is characterised by a short covering rally in the market and possibly by some really “smart” money accumulating at the bottom. The mass psychology could still be summed up by “ok, the market is rallying but it won’t last”. Normally then the market will come down a little bit more (backing up the popular psychology of the rally not lasting) but crucially the market will *not* make a new low. It will then rally again and, once it gets past the peak of the original rally, will *accelerate* as the psychology starts to turn from bearish to “My goodness the market is continuing to rally. I shouldn’t have sold out. Maybe I should think about getting back in again.” Usually at this stage, however, there is a lot of demand from what would normally be considered the professional or institutional investor sector of the market that drives the price higher, and it is usually not until later (near the end of the cycle) that the retail investors come back to join in the fun.

Elliott Wave analysis allows us to identify potential stages of this psychology cycle. The first clue would be a five rally from the bottom price. Of course at that stage we do not know for certain that the price in question will be the ultimate bottom, but we can identify if the market has moved up in five waves. A five-wave move is an impulsive wave and points to the direction of the underlying trend. So if we see a five-wave move up after a downtrend we can raise the probabilities that the trend (at least in the shorter time fractals) is up for the time being and that, after a correction or sideways movement, there should be another move higher. The next clue would be the form of the down move. If, after the five-wave move up there is a three-wave move down, then the odds of another five-wave move up increase dramatically and, because we know that a five-wave move is an impulsive wave, we can be quite confident that the next market phase will be a trending regime. Now, as always the longevity and sustainability of the turn in the market will depend on what time fractal you are analysing. If we were analysing the hourly time frame, then five up and three down from the bottom is useful to know but we will not know whether that price and psychology pattern will turn into something more sustainable over the long term. We simply have to wait and analyse the wave patterns as they evolve over that time fractal.

This point about evolution is important and is at the very core of Elliott Wave analysis. Indeed it is this evolutionary aspect of price action that people dislike about the analysis because the analyst is always at the mercy of “what evolves”. However, in my opinion, this evolutionary aspect simply recognises the fact that a complex system will grow in an evolutionary fashion, and the analyst can never say for certain how that process will evolve. I think most people would agree that it is unrealistic in the extreme to assume that scientists or biologists can predict how the universe is evolving with any degree of certainty. All they have is probabilities and the predictions of future evolution change as time passes and more evidence comes to light. You don’t have to be a religious person to appreciate the sentiments of the *Desiderata* dated 1692 because it aligns itself with an evolutionist’s view of our world too: “And whether or not it is clear to you, no doubt the universe is unfolding as it should.” Things evolve, people and animals evolve. Change is constant.

Why should it be any different in the financial markets? Why do we expect people to give firm forecasts about where the markets will be a year from now, or even a month from now? No one knows that. All we have are probabilities and as those probabilities change on a daily basis, why not accept the fact that the future will take care of itself and focus on the present? As I have previously alluded to, it is human nature to crave certainty and comfort in the face of uncertainty and, of course, this is one reason why something like organised religion is so popular. An evolutionary philosophy though, when it comes to market analysis, is just plain common sense.

In our example of the hourly time fractal and a potential market bottom we would watch to see how the market unfolded and be constantly changing our probabilities as to the future potential of the market. If, in the hourly time fractal the market does 5 up, 3 down, 5 up, 3 down, 5 up, then we now have a five-wave up move in a higher time fractal, and the evolutionary observation increases the probabilities that the *bigger* trend is now up. If this process was repeated in that bigger time fractal then it increases the probabilities that the even bigger trend is up. And so on. Thus the wave analysis is an evolutionary process with an ever-changing probability structure.

By keeping a check on the evolving nature of the wave structure we can begin to see how the fractals align themselves with each other and, in doing that, we will see that the most powerful trending moves happen when there is third wave unfolding as part of another third-wave structure. This “third of a third” is a high probability signal for a trending regime and can be factored into an overall assessment of trading regime analysis. The example of the Aussie dollar in Figure 8.4 shows a good example of the impulsive (trending) nature of a third of a third when the market moved from around 0.5050 in February 2002 to 0.5800 in June 2002, a move of 16% in four months.

THE PART THAT PUTS PEOPLE OFF THE MOST

As mentioned, Elliott Wave analysis has a mixed reputation because generally people either love it or hate it. The people that love it are able to appreciate the fact that

the theoretical basis for it is evolutionary and chaotic, but that there is a hidden order in the chaos. In a very broad sense I would say that these people are generally “right brain thinkers”, meaning that their bias in thinking is driven by a more holistic and intuitive process. On the other hand, the people that hate Elliott Wave analysis are generally “left brain thinkers” whose thoughts are driven by a more logical and objective process, and the one specific area of the analysis that puts people off the most is the area of corrective waves or, in terms of trading regime, the consolidation phases.

A large part of this is due to the fact that while five-wave impulsive trending moves are mostly straightforward and obvious, corrective waves can come in a few forms and are not necessarily too obvious. Nonetheless, if the analyst follows the three rules and three guidelines of the Elliott Wave, then a framework is applicable that will accept or reject certain wave count possibilities.

Dow Theory is one of the pillars of technical market analysis and was developed by Charles Dow (he of the indices) in a series of articles in the early 1900s. One of the main aspects to this theory is that a trend has to be defined by a series of higher highs and higher lows in the case of an uptrend (and lower lows and lower highs in the case of a downtrend). It is only when that series is broken that the analyst can begin to think that the trend is changing direction and so, in that sense, it is recognition of price breaking older turning points that becomes important. Ralph Elliott was heavily influenced by Dow Theory and we can see that influence in his own guidelines concerning wave overlap, where *overlapping waves are generally considered to be part of corrections*. Dow and Elliott theories both have a huge amount in common, but the Elliott Wave goes one stage further to break down and formalise orthodox pattern recognition into its component fractal parts. A symmetrical triangle in Elliott terms consists of certain characteristics which, if they do not appear, would not constitute a symmetrical triangle but rather something else. However, to the orthodox pattern recognition analyst it could still be termed a symmetrical triangle. In a sense these very slight differences between Elliott and orthodox pattern recognition as espoused by Dow, Edwards and Magee are only semantics and it does not alter the analysis from a volatility point of view that corrections or consolidations will generally result in a period of falling volatility.

However, it is this area of corrective waves that the most criticism of Elliott Wave Theory is found because they can be quite complex and, to the layman, look as though they are just being “made up as they go along”. This aspect of Wave theory, that a wave is not found to be labelled correctly until it is completed and even then may change depending on subsequent price action, is undoubtedly the reason why most people give up on the analysis. To reiterate the earlier points, the markets are evolutionary in nature and, in that sense, are continually giving the analyst new information, and this new information (and this is the crucial point) allows the analyst to accept or reject certain outcomes based on the rules and guidelines of the principle. A corrective wave can thus be labelled to have been completed and, therefore, a probability given for future price action; but if that subsequent price action does

not materialise and instead something else happens, then the original probability is rejected with an alternative probability being accepted.

Many people think this is like “having your cake and eating it” but ask an economist to evaluate his anticipation of future events and he will no doubt hedge it by the old chestnut *ceteris paribus*, all other things remaining equal. Well, the fact that Elliott Wave analysis always has an alternative count is exactly the same philosophy. In fact it is better than the economists *ceteris paribus* because the Wave Principle has an alternative in the first place. Most economic analysis would say “such and such should happen, *ceteris paribus*” without giving an alternative if all other things do *not* remain equal!

Corrective waves in Elliott can take a number of forms and this book is not the place to go into the minute detail of such waves but, for the purposes of trading regime analysis, it is worth mentioning a couple of aspects that should help the identification of developing range-trading regime.

OVERLAP

We have already mentioned the fact that five clear impulsive waves should lead the analyst to give a high probability that a correction is due and that the probabilities will be skewed towards this having the characteristics of a range-trading regime. Another way to confirm that a range-trading regime could be developing is to look for wave overlap as this will generally occur in corrective waves. Impulsive waves are trending regimes and will therefore show an environment where pauses in the trend or corrections will be very shallow. Corrective waves, on the other hand, are range-trading regimes and will therefore show an environment where there is no clear up or down trend but rather a sideways trend. This sounds rather simplistic but the point is that if the analyst can identify areas of wave overlap, the simple conclusion should be that a period of consolidation is under way. Sometimes this consolidation can be short but sometimes, depending on the longevity of the previous impulsive move, it can turn into an extended period of range trading giving the analyst who identifies the chop early a distinct advantage.

Let's look at an example of a corrective Elliott Wave and see how it can work out in practice. Again, for a detailed explanation of all the various forms of corrective Elliott Waves I recommend *The Elliott Wave Principle* by Frost and Prechter.

Figure 8.6 shows the bar chart of the USDJPY exchange rate using the daily time fractal from May 2007 until November 2007. This shows the case of the “double three” or extended/running flat and is illustrative of how an initial three-wave correction can turn into an extended range-trading regime. We can see that from the top in June at 124 the market fell in what, by 17 August, was a clear five-wave impulsive movement down to 111.42 (labelled 5). At this point the market corrects in three waves up to 117.14 and the Elliottician would be raising the possibility that this correction was complete with the market ready to fall into new lows. However,




Figure 8.6 USDJPY daily – A–B–C range trade

it is the development of the price action over the next few days that will determine whether this possibility remains just that or whether it is raised to a probability. If it was going to be the case that the impulsive downtrend was set to continue straight away, then the next move down should occur in five clear waves (i.e. an impulsive wave). However, the move down from the high on 23 August at 117.14 to the low on 29 August at 113.86 occurred in only three waves before, crucially, the market then moved back up and *overlapped* with the 115.54 low made on 24 August. This is vital evidence because, by overlapping this low point, the probability that the market is in an impulsive move down has been eliminated. It has been eliminated because one of the rules of Elliott is that wave 4 can never end in the price territory of wave 1 and if the price action from the 117.14 high was to be labelled as a lesser degree (lower fractal) impulsive 1–2–3–4–5, then wave 4 could not have breached the wave 1 low of 115.54.

On 29 August, therefore, the trading regime analyst looking at the Elliott Wave for clues would be weighing up the evidence, and that evidence would strongly suggest that a period of range trading was under way after the three-wave correction from the low on 17 August. The probabilities that the range-trading regime will be extended into the future have risen dramatically by the price action on 29 August. What the price action of 29 August says from an Elliott Wave perspective is that the corrective wave that began on 17 August looks like it will be a *complex* wave consisting of a

series of other corrective waves. This is where most students of Elliott get put off because they expect the analysis to be able to fit into a nicely stylised series of five waves followed by three waves and so on. The existence of complex corrective waves in no way dilutes the theoretical basis of the Elliott Wave but it does, admittedly, make it more difficult to understand.

A THEORETICAL BASIS FOR ELLIOTT WAVES

One of the criticisms of technical analysis from fundamentalists is that it does not appear to have any theoretical basis. “It’s just people drawing lines on charts” they scoff. Granted, the arguments I have made before this chapter, and have dotted throughout this book – that technical market analysis is, at its soul, an analysis of the mass *psychology* of the market participants – have only a theoretical basis that is, if one believes that there is a theoretical basis in psychology as a social science. Undoubtedly there are many people who dispute the validity of psychology as a pseudo-medical field or discipline of any merit and no doubt these people have their reasons for doing so, many of which may well be founded on a fanatical religious faith that dismisses Darwin and his theories on evolution. Nonetheless, the vast majority of people in the world would, I believe, think that human psychology is a subject that is intertwined with who we are and where we came from, and deserves academic status.

However, there is deeper evidence that gives credence to the argument that our decisions as human beings (which have a direct effect on the market price) are driven by natural forces that are at the very core of life on this planet. In the year 1202 a journal called *Liber Abaci* was published (revised in 1228) in Italy by a man who has come to be known as the greatest mathematician of the middle ages: Leonardo Fibonacci. Fibonacci’s fame was such that he was sought after by emperors to provide advice and solve intellectual problems. In 1225 he met Frederick II, who was Emperor of the Holy Roman Empire, the King of Sicily and Jerusalem, among other such titles. The Emperor duly posed a number of problems for Fibonacci and when the revised edition of *Liber Abaci* was published three years later it was dedicated to Frederick.

One of the problems that was posed in *Liber Abaci* is this:

How many pairs of rabbits placed in an enclosed area can be produced in a single year from one pair of rabbits if each pair gives birth to a new pair each month, starting with the second month?

Each pair of rabbits can only come into being after one month, so the number of pairs at the beginning of the first month is 1. Remember, however, that our rabbits are only starting to fornicate at the beginning of the second month of their existence and so, at the beginning of the second month, the number of pairs is also 1. At the start of

the third month, however, the first pair have delivered another pair of rabbits so the number of pairs at the start of the third month is 2. At the start of the fourth month our original pair have given birth to another pair but their first born, given the fact that they cannot begin to produce until the second month of their existence, have not produced. Therefore the number of pairs at the start of the fourth month is 3. At the start of the next month the original pair have produced and their first born (but not their second born) have produced giving the number of pairs as 5. At the start of the next month there would be 8 pairs and then 13 at the start of the month after that. The sequence of pairs begins to take on a logarithmic nature thus:

1, 1, 2, 3, 5, 8, 13, 21, 34, 55, 89, 144, 233, 377, etc.

This sequence of numbers has predictably become known as the Fibonacci Sequence and, because of its basis in natural mathematics, it is one part of mathematics that can be easily understood by non-mathematicians. Perhaps that is why it has become quite popular in recent years, even finding its way into fictional works like Dan Brown’s *The Da Vinci Code*.

It is easy to see from the sequence that the sum of two adjacent numbers gives rise to the next higher number: $1 + 1 = 2$; $1 + 2 = 3$; $2 + 3 = 5$; and so on. After the first four numbers of the sequence the ratio of the two adjacent numbers becomes very interesting.

Figure 8.7 shows the sequence of the two adjacent numbers and the corresponding ratio. We can see that the ratio gradually zones in on a number of 1.618 and the inverse of this number (1 divided by 1.618) is 0.618. Fibonacci had discovered something within this series of numbers that was so simple yet so powerful as to be

Lower number	Higher number	Ratio
3	5	1.6667
5	8	1.6000
8	13	1.6250
13	21	1.6154
21	34	1.6190
34	55	1.6176
55	89	1.6182
89	144	1.6180
144	233	1.6181
233	377	1.6180

Figure 8.7 The ratio of two adjacent numbers

regarded as one of the key balances of life on planet earth. Say hello to the Golden Ratio! This ratio can be expressed as a mathematical constant denoted by the Greek letters Φ or ϕ which stand for phi, and the golden ratio can be found in many creations on planet earth from the natural to the man made. The Egyptian pyramids display the properties of the golden ratio. The Nautilus is a great example of the golden ratio as a spiral. Sunflowers, shells, hurricanes, leaves from trees and human bone structure all have phi at their core. If the universe is fractalised, chaotic and mathematically linked then the golden ratio is critical in this linkage.

Phi links together the fractals and waves we see in the markets, not in a systematic way but in a chaotic fashion with some waves reversing at retracements of, or extensions from, previous waves based on derivations of phi. There is no systematic model to these “Fibonacci levels” acting as support or resistance in the market. Their occurrences are chaotic and therefore cannot be modelled. To left-brain thinkers this, of course, is simply unacceptable. Nevertheless, it is natural mathematics that underpins the Elliott Wave Principle and gives it, in my opinion, a solid foundation. The problem is, of course, that a lot of people do not accept the fact that the golden ratio (or golden mean) is a valid concept in life and certainly not one that is the driving force behind life. For these people the Elliott Wave Principle and Fibonacci numbers when applied to market analysis are anathema, something to ridicule and treat with contempt. As with anything though, it is far easier to ridicule than to attempt to understand. No one is claiming that the markets will move to the tune of the Elliott Wave Principle in some pinpoint accurate Fibonacci-based order all the time. Like orthodox pattern recognition the patterns of the markets will seldom be exactly as they would appear in a stylised description because in the real world we have stop orders, slippage, over- and undershoots to consider. However, if we can accept that there has to be a bit of poetic licence involved in price pattern market analysis (a concept that I agree our quantitative friends will never understand) we can accept the world for what it is and are able to see the patterns in price that are created by market psychology as more of an art form than a science.

It really is amazing to me when I see our more fundamentally oriented colleagues realise that there are many occasions on which markets turn at Fibonacci retracement or projection levels and yet still cannot accept these facts as facts until they can rationalise the explanation for it in their minds. Again it is the *need* for understanding that impedes an acceptance of the reality and the problem when someone asks you why the golden mean has such an influence on the markets is that, unless the questioner accepts the fact that it is the golden mean which is the basic driver of all human life and is the hidden engine of human behaviour, it really is very difficult for that person to accept the thesis. What they are trying to rationalise in their own mind is nothing less than the question of life itself! Why do rabbits breed in a certain pattern? Why are human beings’ joints and dimensions allied to the golden mean? Why were the Pyramids in Egypt built to the specifications of the Fibonacci series? Why is the golden ratio present in the make-up of human DNA and why are the dimensions of the golden spiral manifest in Nautilus, flowers and the spiralling galaxy? These are

questions that certainly cannot be answered by me. However, does this have to mean that an observance of these phenomena and the fact that market prices tend to have the same sort of structure is in some way a kind of leap of faith in some sort of alchemy? I do not think so, but there are many out there who do. Maybe we should turn the question on the fundamentalists and ask them why a market price should act in a certain way when a piece of information reaches it. Why should the stock market go up when GDP is going up? Why should it not go down, because rising GDP might stoke inflation and cause interest rates to go up, thus lowering the return on the stock market? Unfortunately we have become conditioned to thinking by rote about the economy and markets in the last half-century or so, and we cannot accept any alternative theory about market behaviour unless it can be fully rationalised in those terms. Is it not about time we opened our minds and thought outside the CAPM and macro-economic box a little more?

SUMMARY

In this chapter we have examined how the Elliott Wave Principle can present the trading regime, or volatility, analyst with situations where the probability of increasing or decreasing volatility (range expansion and contraction) is relatively high. We have seen how impulsive moves in the market differ from corrective moves and how this interaction is a natural phenomenon steeped in the natural laws of mathematics. Human nature dictates the ebb and flow of market direction and the cycle of fractal-based evolving complex systems is such that price expansion is followed by price contraction and vice versa.

The Elliott Wave Principle is built on sound human behavioural, psychological and Chaos Theory. It explains the personalities of market cycles and how price action anticipates changes in the so-called fundamentals. It is a trend-following, range-trading, money management plan all rolled into one. It is, in my opinion, the most complete investment and trading methodology that exists.

9

Moving Average Envelopes

The humble moving average is probably one of the most useful indicators in market analysis. Moving averages can determine trend direction but they can also be used as an overbought, oversold oscillator if we measure extreme readings away from them. They are the backbone of technical market analysis and appear in some form or another in most methods and indicators. Look, even economists use them! I always like to rib economists that they are actually trend followers in disguise because what they tend to do is to look at a chart of some economic data – for example, retail sales – and, in order to filter out the “noise” in the series, they put a moving average on it to determine the underlying trend. This is what economists do. They try to determine the *trends* in the economy in order to predict the future course of market price action (however disputable that may be) but, in order to determine the trend direction, they look at a moving average of an economic data series and then assume that this trend will continue into the future! Exactly as trend followers of market price data do. In methodology there is little difference between an economist and a trend follower, but in ideology there is, of course, a massive gap as the economist assumes that by identifying the trend of some economic variables – which may or may not be correlated to the market price – he can determine the future course of market prices, whereas the trend follower identifies the trend of the market price itself and follows that and that alone. Therefore, the moving average can be thought of as one of the pillars of all market analyses. Ironically enough there is a very valid argument for saying that economists should not really be looking at moving averages on their macro-economic data but should instead be using oscillators to identify when something is overbought or oversold, because the long-term macro-economic data will tend to move in cycles (the business cycle), and although moving averages will pick up a lot of the trend from overbought to oversold and vice versa they will lag many of the inevitable turns in the cycle. Nevertheless, the moving average remains the most popular of all statistical techniques.

Besides looking at moving averages for trend direction and trend extremities we can also look at moving averages to help us to identify the market's trading regime. This is done through a technique known as moving average envelopes (MAEs).

The concept of moving average envelopes is simple. All we do is take a moving average of the market price and plot a line a fixed percentage above the moving average and plot another line a fixed percentage below the moving average. In this sense we have, around a moving average, created bands that form an "envelope" around the market price action.

Figure 9.1 shows the bar chart of the EURUSD exchange rate in the daily time fractal with moving average envelopes. The dashed line is the 21-day moving average and the solid lines are plotted at 1% above and 1% below this moving average. We can see that sometimes the price action trades inside these envelopes and sometimes the price action trades outside these envelopes.




Figure 9.1 EURUSD daily – moving average envelopes

This is where it gets interesting. Most people think of moving average envelopes as indicators that will point out support and resistance areas. The theory is that the lower envelope should provide support for the market price and the upper envelope should provide resistance for the market price. The logic of this is flawed, I believe, because it assumes that a market's price action will be constrained by some fixed percentage above and below a moving average. Now, I do believe that analysing the distance that the actual market price is away from a moving average can provide a good overbought or oversold oscillator indicator similar to the analogy of stretching a rubber band – at some point it has to get to the point of maximum stretch. However, the notion that a market's price action will be constrained by areas at a fixed

percentage away from a moving average assumes that the market in question will be going sideways most of the time, and this is clearly not the case in most markets. We can see in Figure 9.1 that there are times when the market price clearly trades above the upper band for a long period of time.

I prefer to use moving average envelopes as trading regime indicators by examining when the market is trading inside or outside the bands. If the market is trading inside the bands, then we can say that the price action is being contained at that time within a band of prices above and below the moving average. Therefore, in all probability this will result in a choppy, range-trading type of price action. However, if the market is trading outside the bands then we can say that the price action is not being contained within a band of prices above and below the moving average and therefore, in all probability, the market price action is behaving in a strong trending fashion. In this way the moving average envelopes become an indicator of trading regime and can be incorporated into an overall analysis of trading conditions.

Now there are certain issues with this indicator that have to be addressed. The first is the usual issue with regard to what length of moving average to use. As I have mentioned previously, I am not a believer in optimising variables like moving averages but instead look to select a moving average on the basis of common sense of what I am trying to achieve. In using moving average envelopes the goal is to determine when the market is in a trending phase and when the market is in a range-trading phase; therefore, in order for me, as a trader, to benefit from the analysis, I do not want to have too short a moving average that will be chopped around with the market price data, nor do I want too long a moving average that will hardly move with the market price data. I want something in between that will capture the market's price cycles without being too sensitive or insensitive. By selecting a moving average on the basis of common sense I might not be selecting the moving average that has the best "fit" to past data but I can rest assured that my logic of selection is sound.

The second issue to address is the level at which we want the fixed percentages to be above and below the moving average. Again, common sense has to be our guiding light and it generally boils down to the time fractal the analyst is dealing with. If the analyst wants to analyse trading regimes at the micro level or the very short term, then because price fluctuations will be lower than they would be in the longer term a smaller fixed percentage should be used. Conversely, if the analyst wants to analyse trading regimes at the macro level or the very long term, then because price fluctuations will be larger than they would be in the shorter term a higher fixed percentage should be used.

Figure 9.2 shows the bar chart of the EURUSD exchange rate in the weekly time fractal and moving average envelopes of 1% either side of a 21-week moving average. We can see instantly when comparing this chart to Figure 9.1 – which shows the same moving average envelope variables but in the daily time fractal – that the band on the weekly chart appears much narrower than on the daily chart. This is because the price data in the weekly chart will have greater fluctuations (bigger bars)



Figure 9.2 EURUSD weekly – moving average envelopes

than those in the daily chart and, of course, this is a function of the time fractals. A monthly chart would show even greater fluctuations.

Therefore, the trading regime analyst has to consider the time fractal and change the fixed percentages of the envelopes accordingly. For longer time fractals the analyst will want to use bigger percentages to find sensible cycles in the market with respect to regime, and for shorter time fractals the analyst will want to use smaller percentages to find sensible cycles in the market with respect to regime. The analysts will also want to consider the type of market he is dealing with. As individual stocks will fluctuate in percentage terms much more than currencies, the former will generally have bigger fixed percentage bands than the latter.

The reader will observe that there are obviously times when the market price will blip above the bands only to return to trading inside the bands, and there will also be times when the market price can trend in a certain fashion while still being contained within the bands. Nothing is going to be perfect, however, so we shouldn't get too bothered about this. Remember, what we are doing as trading regime analysts is looking for situations where the *probabilities* favour a trend or where the *probabilities* favour a range. In all market analysis we deal with probabilities and we never deal with certainties, because there are none. Chapter 22 deals with how best to use trading regime analysis, and that discussion deals with combining our trading regime indicators in order to maximise this subjective analysis of probability.

Moving average envelopes allow us to see what trading regime the market is currently in and so it is not predictive. However, as with most, if not all, market analysis the inference is that there will be a degree of persistence involved after an observation is made. Borrowing again from our friend Keynes, the evidence is presented to the analyst of market price or trading regime and, until the facts change, the assumption must be that the recent past will continue. If evidence is presented that disputes

at a distance above and below. Keltner's centre-line is a 10-period simple moving average of the "typical price", which is the average of the high, low and close.

The upper and lower lines are calculated by the distance away from the centre-line of the simple moving average of the past 10 periods' trading ranges (high to low). Keltner's use of the indicator was to regard a close above the upper line as strongly bullish and a close below the lower line as strongly bearish, so in this regard he was using it as a non-continuous trend-following, volatility breakout type of system. He did also stress, like most good analysts should, that other indicators should be used to confirm the breakout.

There have actually been many variations of "channel" type indicators and systems over the decades, but the key thing they all have in common is that trading within the bands is considered to be noise/range trading and trading outside the bands is considered to be breakouts/trends.

SUMMARY

This chapter has focused on using the humble moving average as a trading regime indicator by plotting shadow lines at a certain percentage above and below the moving average. We have used the methodology in a standard trend-following manner in that we use the breaking of the envelopes as signals that the market price could be about to enter a strong trending regime and only when the price is in between the envelopes do we assume that a range-trading regime is in existence. This methodology means that only when the market is in a strong trend will the regime be identified as a trending regime, and if the market is trending in a weak manner it will *not* be identified as a trending regime. In this sense it can be thought of as a channel or volatility breakout method of analysis.

Using price envelopes or bands at a fixed percentage around a moving average allow the trading regime analyst to define when a trend is strong or weak and, hence, define the probability of the trading regime. Bands around a moving average, however, need not be fixed and can be variable. In the next chapter we will explore how these variable bands can add a slightly different dimension to volatility analysis.

this (a break above the moving average envelopes for example) then the facts of the case have changed and the “view” of the analyst should change too.

Figure 9.3 shows the bar chart of the Nasdaq Composite Index and the associated moving average envelopes from 1995 until 2007 using the monthly time fractal. We can see that the envelopes do a good job of defining when the market is in a strong trend, either up or down, and when it is in a weak, sideways, choppy trend such as in the move up from 2004. The market was making higher highs and higher lows during that period and so the strict definition of an uptrend was there, but the trend was incredibly weak and choppy. Trend-following systems would have really struggled to make any sort of return during this period but, if the risk capital allocated to such systems was altered to reflect the fact that the moving average envelope analysis suggested that a range-trading regime was most likely in force, the draw downs could, in the very least, have been minimised. In fact, capital could also have been allocated towards mean-reverting type systems and that would have enhanced returns during this period.




Figure 9.3 Nasdaq Composite Index – moving average envelopes

VARIATIONS

There are many variations of the moving average envelope technique and one of them came from Chicago grain trader Chester Keltner (1909–1998). He was a commodity trader who, in 1960, wrote a book called *How To Make Money In Commodities* in which he described “the ten day moving average trading rule” that subsequently became known as Keltner Channels. It is, as the name suggests, a variation of the moving average envelope showing a central moving average line plus channel lines

10

Bollinger Band Width

We have just seen how moving average envelopes can help us in our trading regime analysis by plotting lines at a certain percentage above and below a moving average of the market price. Now imagine that, rather than plotting the lines at a certain fixed percentage above and below the market price, we plot the lines at a variable percentage above and below the market price, depending on the underlying volatility of the market. Technical analysis legend John Bollinger has created an indicator that does exactly this, and it has become known, unsurprisingly, as Bollinger Bands.

Bollinger Bands have become a very widely used and useful indicator for market analysis because they take the concept of moving average envelopes one stage further to take account of changes in market behaviour and volatility. Bollinger realised that it is the volatility of the underlying market that is vital to an analysis of the potential range of prices that the market may explore, and that the volatility of markets, whether it be indices, individual stocks, bonds, commodities or currencies, can and will differ, sometimes a great deal. Moreover, he realised that this volatility will change over time depending on the underlying market conditions, which themselves will be affected by the psychology of the market participants. Therefore, having an indicator of limit bands around price action that was dynamic and changed according to the ebb and flow of the markets was a terrific addition to the body of knowledge in technical market analysis.

The difference between moving average envelopes and Bollinger Bands is that instead of plotting the bands at a certain fixed percentage above or below a moving average of the price, we plot the bands at a specified number of standard deviations above or below a moving average of the price (the default number of standard deviations that Bollinger used is two). What this means is that the bands will fluctuate in width as the volatility as measured by the dispersion of prices fluctuates. It is very important to note that we are measuring the standard deviation of the population itself and not the standard deviation of the changes in values, as we do in measuring historical volatility. This means that the Bollinger Bands indicator, although a

measure of market volatility, is not the same as measuring the historical volatility of the market.

The calculations for Bollinger Bands are straightforward.

SMA = the simple moving average of the closing price over the last N periods.

SD = the population standard deviation of the closing price over the last N periods.

Upper band = $SMA + (SD * 2)$

Lower band = $SMA - (SD * 2)$

Figure 10.1 shows the bar chart of the EURUSD exchange rate in the weekly time fractal along with the 20-period moving average (dashed line) and the associated Bollinger Bands (solid lines). We can see that the market price (the bar chart) is mostly contained within the Bollinger Bands and that, when it strays outside the bands, it doesn't usually stay there for too long. However, note that this is not the same as saying that we should sell the market if the price goes above the upper band or buy the market if the price goes below the lower band. Yes, there can be times when the market price strays outside the bands only to reverse direction quite sharply but there are also times when the market price strays outside the bands only to continue firmly in the same direction, albeit from subsequently inside the bands.




Figure 10.1 EURUSD weekly – Bollinger Bands

As trading regime analysts, however, we are not as interested in the potential direction of prices as we are in the potential for a strong trend, either up or down, in price or a range-trading period in the price, and the Bollinger Bands give us one of the best trading regime indicators that we can use in this regard. Notice from the chart that the difference between the upper and lower Bollinger Bands is variable, with it being sometimes wide and sometimes narrow. This is the mathematical certainty of the calculation because we are plotting the standard deviation of the price above

and below a moving average of that price and, of course, as the standard deviation is a measure of dispersion it will sometimes be wide and will sometimes be narrow depending on how the market is behaving. Therefore, the varying *width* of the Bollinger Bands should be a useful indicator of the volatility of the market as defined by dispersion of the price series.

In fact, John Bollinger defined the Bollinger Band Width separately and it is generally shown as a separate chart along with the chart showing the bands around the price on a time series. The Bollinger Band Width is defined as:

$$(\text{Upper band} - \text{Lower band}) / \text{Middle band}$$

where the middle band is the moving average of the price.

Figure 10.2 shows the same EURUSD chart as Figure 10.1 but this time with the Bollinger Band Width shown in the bottom chart. When it is plotted in this way it is much easier to observe the variations in the width of the band and how the width of the band reaches an extreme on the wide side and on the narrow side before reversing. What becomes clear is that the market price tends to trade within a range when the




Figure 10.2 EURUSD weekly – Bollinger Bands and Band Width

the bottom chart. We can see that during late 2003 and early 2004 the share price of Google was trading in quite a narrow range below \$200 and that, as this was happening, the Bollinger Band Width was falling from around 0.9 to around 0.2. However, when the market price broke out of this range and started to rise to over \$300 (a 50% rally!) it coincided with a rise in the Bollinger Band Width back up to the 0.8 area. Thereafter, another period of range trading occurred which brought the band width down again before yet another trending period began, this time from \$300 to \$500, which also coincided with a rise in band width. We can see from the chart that whenever there is a period of range trading the band width falls, and when there is a trending period in the price, the band width rises. What is really useful about the Bollinger Band Width analysis, however, is that it becomes very easy to identify periods when the market is vulnerable to a breakout in price action after being in a period of range trading. In the Google example we can see that whenever the Bollinger Band Width falls to below 0.2, it is followed by a period of trending or what can also be referred to as a “volatility breakout”.

Volatility breakout is a generic expression in the markets used to describe a situation whereby the market price “breaks out” from a period of low volatility (generally associated with a range-trading regime) to a period of higher volatility (generally associated with a trending regime). Volatility breakouts are arguably the best way to trade the markets because of the high probability they have of identifying a relatively large trending move in the market price, but it very much depends on how the analyst measures volatility that makes the difference to whether this strategy is seen as successful or not. To stress yet again, here we are measuring volatility by the dispersion of a population of prices (the compression or expansion of price ranges) and not by the dispersion of relative changes in the market price. There is a big difference and this difference is perhaps one reason why some people say, “volatility breakouts do not work”. I fully agree that, if you are measuring volatility by the historical volatility standard then not all low volatility periods will be followed with a trending move in the market because the volatility could be falling when the market price is actually trending. However, when we use the standard deviation of the price *series* and not the *changes* in price as our volatility indicator, then it becomes a mathematical certainty that a trending period will be met with a rise in the volatility indicator, and therefore a period of low volatility gives a great opportunity for the analyst or trader to benefit from the inevitable subsequent trend or price range expansion.

Bollinger Band Width is therefore one of the best indicators that the analyst has at his disposal to spot potential volatility breakouts into trending regimes and the market psychology associated with the volatility breakout makes it a trading strategy with a great deal of potential.

The question, of course, that readers will be asking themselves is: What particular levels of Bollinger Band Width should we define as “low” or “high”? This question is of paramount importance to the quantitative analysts and model traders among us, but it becomes of less importance when we want to use the Bollinger Band Width as a discretionary tool for evaluating trading regimes. For our quantitative friends, I

band width is falling and it tends to exhibit trend-like qualities when it is rising. This is what we would expect because if the market price is trading within a range, then the standard deviation of the price series *should* be falling; and if the market price is trending, then the standard deviation of the price series *should* be rising. It is very important to reiterate that the use of these population standard deviation-based measures of volatility will give us a much better reading of range trading and trends than the use of the standard deviation of the changes in the price, which is the historical volatility calculation. As we have seen in Chapter 3, a market can be in a strong trend, but the historical volatility measure is actually falling; or it could be trading in a narrow range, but the historical volatility measure is actually rising. However, like the plain vanilla standard deviation, the Bollinger Band Width indicator is very useful for identifying when a market is trending because the indicator will generally rise until the trend starts to dissipate and it is also very useful for identifying when a market is range trading because it will generally fall until the market price breaks out of that range.

Figure 10.3 shows the bar chart of the share price of Google, the internet company, in the weekly time fractal with the associated Bollinger Bands and Band Width on




Figure 10.3 Google weekly – Bollinger Bands and Band Width

am afraid to say that there is no generic answer to this question because each market will be different, and even if we standardised the scale each market will still exhibit different behavioural tendencies, such that one market might give low band width readings at different levels from another. It is easy to rank the width readings in a spreadsheet that gives the analyst a gauge of where the current reading is in relation to past history, but the possibility of new highs or lows in the width still, of course, exists.

Nonetheless, I do not believe that this dilutes the potency that the Bollinger Band Width has as a trading regime indicator, because the general logic still has a lot of power. A big market move will generally come after a period of not much happening, and a period of not much happening will come after a big market move. This is simply the way markets work because it is rooted in the behavioural cycle of human beings, and if we have an indicator that clearly shows us the periods when not much is happening as well as the periods when a big market move is happening, then we can use this to determine our overall game plan for the market.

Bollinger Band Width clearly shows us the direction of the standard deviation of a price series even if it does not have an upper limit, or even if it does not always go down to a certain low on every occasion. If we use it as a discretionary tool for evaluating trading regimes, then it can add value. This will be anathema to the quantitative community but so be it. Our job, as I see it, as trading regime analysts is to identify the discretionary probabilities of whether a market will be trading in a range or trading in a trend and, although many models can be created that attempt to mathematically predict volatility, it is not necessarily the same thing. We can certainly make use of such mathematical models and tools but my preference is to use such models as part of an overall process encompassing a number of techniques.

Figure 10.4 shows a bar chart of the logarithmic change in the price of the copper future in the monthly time fractal with the associated Bollinger Band Width. We can see clearly that when the Bollinger Band Width is falling then the market is generally range trading, and when the band width is rising then the market is generally trending. Notice again how the band width tends to have a lower limit at which, when it is near to that level, the market eventually makes a trending type move and the band width subsequently widens again. When copper started rallying in 2003 the band width exploded higher from its lows but then trailed back down again as the market consolidated in a range-trading regime in 2004 and early 2005. However, the band width did not come fully back to its historical lows of around 0.2 before the market then took off again in a huge bull run, and this is undoubtedly one aspect of using the Bollinger Band Width as a trading regime indicator that is not entirely satisfactory. Of course, in a utopian world the markets would behave in such a manner that the band width always came back to the historic low before taking off in a big trending move again. Unfortunately no market will behave in such a manner all the time so we, as trading regime analysts, need to take a holistic approach in coming to our conclusions, and that means combining our analysis techniques in a more discretionary approach.

good old eyeball method. As one legendary technical analyst is reputed to have said once when trying to get the message across, “Technical analysis does not have to be complicated. Look at the line on the chart. If it is going up that is good. If it is going down that is bad” and I would very much agree with those sentiments.

If we used an analysis of Bollinger Band Width as part of an overall discretionary process we could see instantly whether the indicator is moving up or moving down. Of course, we could get a computer to do this for us but the lags involved in measurement would sometimes work against us. When it comes to observing a measure of the direction of volatility in order to work out the potential probability of the market’s probable trading regime, a case can be made for just looking at the “line on the chart”. So, in the case of copper above, the fact that the Bollinger Band Width had started to rise again by the end of 2005 would have alerted the analyst to the observation that, although volatility was not at historically low levels, volatility (as measured by range expansion) was rising again. Therefore, the trading regime should be regarded as trending.

SUMMARY

In this chapter we have seen how an analysis of the width of Bollinger Bands can be used to determine the current and likely trading regime of the market in question. Standard deviation is an excellent measure of range expansion or contraction and combining standard deviation analysis with a moving average in this way can give us a very visual representation of when the price is compressed in a range and ready to break out into a trend. It can also give the analyst a good idea of when a trend is in an extreme position and, therefore, when the probability of trend exhaustion will set in.

In the next chapter we will look at another classic indicator that can be used to determine the strength of the trend in such a way as to give the analyst a probability in their mind of the market’s probable trading regime.

11

The ADX

J. Welles Wilder Jr is a name synonymous with technical analysis and he is best known for the technical indicators he developed, which are now core requirements in any technical analysis software. Welles Wilder was a mechanical engineer and a real estate developer before concentrating his efforts full time on the markets, where he went on to develop, among others, the Relative Strength Index (RSI) which is one of the most widely used and well-known oscillators around, and also the Average True Range (ATR) which, we will discuss later but is used in many ways, not least in the determinant of position sizing in a very famous trading system. However, another indicator that he developed has become a standard in identifying trend strength, and is therefore very relevant to a discussion about trading regime analysis. That indicator is the Directional Movement Indicator.

Just as the name suggests, the Directional Movement Indicator measures the directional movement of the market price. The logic behind the indicator is that it assumes that directional movement can be considered strong when the market price is making higher highs in the case of an uptrend and lower lows in the case of a downtrend. Therefore the indicator assumes that the current period's high should be higher than the previous period's high in the case of an uptrend and that the current period's low should be lower than the previous period's low in the case of a downtrend.

The Directional Movement Indicator calculation is based on the concept of directional movement (which can be both positive and negative) and is calculated using the following rules:

1. If the current period's high price is greater than the previous period's high price, and if the current period's low price is greater than or equal to the previous period's low price, then the directional movement is positive:

$$\begin{aligned} +DM &= \text{High price (current period)} - \text{High price (previous period)} \\ -DM &= 0 \end{aligned}$$

2. If the current period's low price is less than the previous period's low price and if the current period's high price is less than or equal to the previous period's high price, then the directional movement is negative:

$$+DM = 0$$

$$-DM = \text{abs}(\text{Low price (Current period)} - \text{Low price (Previous period)})$$

3. If the current period's high price is greater than the previous period's high price, and the current period's low price is less than the previous period's low price, then the directional movement is positive if the change in the high price is greater than the change in the low price, and the directional movement is assigned as per rule number 1. However, the directional movement is negative if the change in the low price is greater than the change in the high price, and the directional movement is assigned as per rule 2.
4. For all other permutations

$$+DM = 0$$

$$-DM = 0$$

These rules set out the “directional movement”. The other variable that goes into the calculation of the indicator is something called the True Range. The True Range is defined as the largest of the following:

- (a) The difference between the current period's high and the current period's low.
- (b) The absolute difference between the current period's high and the previous period's close.
- (c) The absolute difference between the current period's low and the previous period's close.

When this True Range has been established it can be averaged over a time series, and this then creates the Average True Range or ATR which is, as mentioned, one of Welles Wilder's most valuable indicators. For more on the Average True Range see Chapter 15 on Donchian Channels.

Once we have the Directional Movement (DM) and the True Range (TR) we can then calculate the Directional Movement Indicator using the following formulae:

$$+DMI = \frac{\text{Average of } +DM \text{ over previous } N \text{ periods}}{\text{Average of TR over previous } N \text{ periods}}$$

$$-DMI = \frac{\text{Average of } -DM \text{ over previous } N \text{ periods}}{\text{Average of TR over previous } N \text{ periods}}$$

where N = the period over which DMI is being analysed (Welles Wilder's recommendation was to default to 14).




Figure 10.4 Copper monthly and Bollinger Band Width

BOLLINGER BAND WIDTH DIRECTION

We have seen how an analysis of the level of the band width is important in determining when the trading regime is potentially overstretched. High band width eventually leads to trend exhaustion and low band width eventually leads to volatility breakouts. However, if the analyst just looks at the direction of the Bollinger Band Width then he can get a good idea of where the balance of risk lies in terms of trading regime. Once more I am going to incur the wrath of the quantitative community by saying that it is useful to analyse the direction of the band width from a discretionary basis (although I would refer more objective minded analysts to the Trading Regime Indicator in Chapter 21, where I look at standard deviation in relation to a moving average).

There are a number of ways to look at a time series to determine the trend of that series. We can put a moving average on it and determine whether the series is above or below the average. This has the advantage of smoothing out the blips up and down of the time series itself, but it also has the perennial disadvantage of being a lagging method. Quick changes in a time series will not be captured by a moving average or a lagging indicator. There is no way round this problem. The quants can optimise and data fit all they like but the basic fact remains that a moving average type of indicator will always lag the underlying time series. This is not to say that it does not have its uses and it can definitely add value as long as we are aware of its limitations. An alternative method, of course, in analysing the trend of a time series is to use the

So far so good. We now have a measure of the directional movement of the market and these measures are used as part of the indicator to indicate buy and sell signals based on when the +DM and the -DM lines cross each other. This is fine for analysts looking for market direction but what we, as trading regime analysts, are really interested in is how we can use this data to determine the strength or lack of strength of the directional movement. Welles Wilder provides us with such an analysis by way of the Average Directional Movement Index or ADX.

The ADX measures the strength of the trend regardless of whether the trend is positive or negative. It is calculated as:

$$\text{ADX} = \text{Average of DX over the previous } N \text{ periods}$$

where

DX = the total range of directional movement

$$= 100 * \text{abs}(+DMI) - (-DMI) / ((+DMI) + (-DMI))$$

The way to interpret the ADX is that the higher the value of the ADX the stronger the trend, and the smaller the value of the ADX the weaker the trend. Welles Wilder recommends that an ADX value above 25 indicates a market that is strongly trending, and an ADX value below 20 indicates a market that is not trending. So here we have an indicator that sets out to determine the strength of the trend of a market and it is therefore defining when a market is trending (in a trending regime) or not trending (in a range-trading regime).

When the subject of trading regime analysis is discussed most people who have used the DMI and the ADX point to this as the benchmark for analysing whether a market is trending or not. It is probably the most widely known indicator of the strength of a trend. However, it is worth examining how it behaves in practice for the purposes of inclusion in our overall analysis set.

Figure 11.1 shows the bar chart of the US 10-Year Treasury Bond Yield in the monthly time fractal and the associated 14-period ADX chart below it. We can see that in 1989 the Average Directional Index was still around 25 but was falling. The fact that the ADX was falling should have alerted the trading regime analyst to raise the probability that the bond yield was likely to be trading in range-trading fashion. In 1990 there were a couple of attempts made by the ADX to stabilise and start rallying but the index could not manage more than a couple of months up move before falling back again. However in October 1991 when the bond yield was trading at 7.43% the ADX had managed to have three up months from a very low bottom (of 11 in August). At this point the trading regime analyst would be thinking that the ADX, at 11, had been so extremely low that it could not get much lower and that it now seems to be putting in a sustainable rally. It would appear to be a good time to switch the bias of trading regime from range trading to trending. So it proved, with the bond yield falling in a trend-like fashion from 7.43% in October 1991 to a low of 5.16% in October 1993. At this point in October the ADX reading attained very high historical levels of around 40.

that the Average Directional Index moved down towards the historically low levels that were seen in 1991.

At this point, in mid-1997, the analyst should be asking the question of how much further down the ADX could go, and when, in July 1997, it started to turn back up, the bias should have been switching back from range trading to giving a trending regime more probability. At this point the US 10-year Treasury Bond was trading with a yield of just over 6% and over the next 13 months that yield fell to a low of 4.16% in an impulsive trend-like fashion.

I must continually emphasise the point that we are dealing with probabilities in everything we do in the markets. While I recognise that the more mathematical and quantitatively focused among us would point out that the blips up in the ADX in, for example, 1994 and 1995, would be grounds for switching to a trending bias, I would argue that because the ADX was not trading at really low levels I would only have put a mild bias on a trending regime and not a strong bias. So, yes, we must pay attention to the direction of the ADX but we can alter the strength we put on those signals by shaving the probabilities, depending on where the index is trading on its potential historical scale.

SUMMARY

This chapter has examined a classic indicator of trend strength and found that it can be a useful indicator of the current and future trading regime of a market. By looking at the direction and not just the level of the ADX, the analyst can add this evidence to his overall trading regime analysis.

12

Point and Figure Charts

“Wall Street never changes. The pockets change, the suckers change, the stocks change, but Wall Street never changes because human nature never changes.”

Jesse Livermore

One of the most useful techniques for charting the markets is also one of the oldest. Before modern-day charting techniques became much easier to generate, thanks to the help of the computer, market participants were forced to either mentally or physically record the price action themselves. People who mentally recorded the price action became known as “tape readers” as they would watch the ticker tape recording of the stock market (which showed the price and volume going through the market) and they would mentally be able to record this information in their heads in order to help with their decisions. The old version of the ticker tape was a machine that spat out the prices in a paper tape form and the modern version is the horizontal electronic screen that we are all familiar with. For example, the tape readers might be watching the ticker tape and observe that a stock had rallied up to a level of \$20 but could not seem to break through that level. They might therefore conclude that there was no more demand left in the market and the stock would probably fall back, so they would go short the stock. They would have recorded and assimilated this price information in their heads alone without physically recording the data.

One of the most famous tape readers of all time, and arguably the greatest trader who ever lived, was a man in the USA called Jesse Livermore. Livermore started trading the markets when he was a teenager in the late 1800s in what were known as bucket shops. These were places where people could go and trade stocks on margin, were generally a law unto themselves and were not unlike bookies or betting shops. Livermore became so adept and successful at trading the markets through these places that he became known as “the boy plunger” and was banned from a few of them because, as the owner of the place acted a lot like a bookie, when Livermore won, he lost. Livermore went on to become one of the wealthiest men in America,




Figure 11.1 US 10-Year Treasury Yield and ADX

The month of that low (October 1993) was a distinct hammer bottom on the Japanese candlestick chart and the month after that low the Average Directional Index started to fall. So at this point the trading regime analyst should be thinking of switching the bias away from trending (thinking that the downtrend has potential exhaustion) and back to a bias for range trading. As it turned out, the move back up in yield proved to be quite a powerful one as can sometimes happen from trend exhaustion points, and the yield moved all the way back up to 8% by November 1994. At that point the ADX actually started to move up again giving the trading regime analyst pause to consider the potential for another trending regime to be the favoured bias, but when, in February 1995, the ADX dropped again as the yield moved back down to the mid-7% area, the analyst would be thinking of range trading again. Don't forget that the ADX was still at very high levels (above 30) and so the balance of probability lay in the fact that there was still some consolidation to come.

There was another small blip up in the ADX in late 1995 as the yield fell to 5.5% but again the market reversed relatively quickly, and by February 1996 the ADX was dropping again. By this time the big range-trading regime was in full force with the market starting to coil quite markedly over the next year and this price action meant

and made and lost a fortune a few times (which qualifies him as being regarded as one of the best, if not the most consistent, traders of all time). He had his own train, lived like a king in Florida and in the crash of 1907, when Livermore was 31 years old, JP Morgan himself pleaded with Livermore to lift some of his short positions in order to support the market and stop the banking system from going bust. In 1940 however, Livermore took the “ultimate stop loss” when he shot himself in the head, having suffered all his life with periods of clinical depression that had been exacerbated with family troubles during the 1930s. One of the all-time classic books on investment, *Reminiscences of a Stock Operator*, is based on Livermore’s life and is generally regarded as required reading for anyone interested in trading and trading psychology.

Livermore had this big office in New York where he would employ clerks as “markers” to record the prices of stocks on a massive blackboard on a live basis and he would stand and watch this price action before deciding to “take a plunge” into a stock or to sell it short. This was a progression from mentally recording the price action to enable him to still physically see the most recent prices as well as the current prices, and the market in general had progressed to participants manually recording the price action themselves in this way.

In the early 1900s traders on the floor of the stock exchanges were recording price action manually on a sheet of paper or in a notebook. However, rather than record every single price change that went through a stock, which would be an impossibly time-consuming task, they started recording the general movement of prices by first of all deciding what constituted the minimum price movement they would record and then mark an up move of this minimum price with an X and a down move of this minimum price with a decimal point. Thus was born the Point (.) and Figure (X) charting technique. The Point and Figure (P&F) technique has proved to be one of the most useful techniques for recording price action, but especially so with regard to identifying breakouts in price action from a range as well as for identifying that range itself. It is therefore a useful technique for the trading regime analyst to have in his armoury.

The decimal point was eventually replace by “O”, so the modern P&F chart looks like a series of Os and Xs which, of course, the fundamentalists, who have never bothered studying the technique, deride as some sort of infantile game of noughts and crosses! I gave up worrying about these people a long time ago when I realised that some people will just never be pragmatic enough to maintain an open mind, so why bother trying to convince them of the validity of a trading or investment technique when they won’t even bother to listen. A lot of people in the market now “play noughts and crosses”, but this “game” is generally a profitable one for those who play it. Let’s look at an example of a P&F chart.

Figure 12.1 shows a Point and Figure chart for copper, and the first thing to notice about this chart is that the horizontal axis has no time scale with only the vertical axis having a price scale. We can see that there are a series of columns with Xs representing upward price movement and Os representing downward price movement.




Figure 12.1 Copper – Point and Figure 10 × 3

This is a 10 × 3 chart, which refers to the box size and reversal amount. The box size refers to the price movement that the analyst wants each X or O to represent which, in this case, is \$10. The reversal amount refers to the amount that the price has to reverse in order for a new column to be created on the chart, and in this case it is 3. What this means is that the price has to reverse by \$30 (10 times 3) in order for a new column of either Xs or Os to be created. So when we see that one column ends and another has been created, this means that the price has reversed (up or down) by \$30.

One of the best features of P&F charts, and the one that we are especially interested in as trading regime analysts, is the clear way in which support and resistance levels appear. In Figure 12.1 we can see that the price of copper had a lot of difficulty getting up past the \$154 to \$164 zone and found support at around the \$114 zone as well as the \$84 to \$74 zone. Due to the fact that there is no time scale on the chart, an analyst with no perception of where copper has traded historically would not necessarily think that there was anything special about this range-trading regime, but the fact that we are using a \$10 box size for a market which, at that time, was

trading around the \$100 level should tell the analyst that this is analysing the long-term price movement. The longer the time fractal that we want to analyse in P&F charts, the bigger our box size should be; and the shorter the time fractal we want to analyse, the smaller our box size should be. In fact the range trading that occurred in copper between around \$70 and \$160 lasted for about 17 years and, although P&F analysis does not include a time element and only concentrates on pure price, if we are using a big box size like this and see a period of range trading like this, we should pay very close attention because we know it will have occurred over a long time frame.

Moreover, being aware of the time frame gives us an extra edge when it comes to analysing a breakout from that period of range trading. In this case, when the price of copper breaks up past the \$154 to \$164 resistance zone the analyst will know that this is a rare event, and should be prepared for the likelihood of a strong trending regime to develop. Of course there is the possibility that the price does not enter a strong trending regime and just goes on to range trade at different zones, or falls back into the old range-trading zone, but being aware of the fact that the price is breaking up past these resistance levels for the first time in 17 years of historical price action adds a lot of weight to the *probability* that what is about to happen might be out of the ordinary.

And so it proved, with copper entering into one of the strongest, multi-year trends ever seen in a commodity market. The price moved from the breakout level of around \$170 to a high of over \$410 without ever having a pullback in price of \$30 or more, and this price movement is represented by the very tall column of Xs on the chart. Using this analysis, the trading regime analyst would have put the probability of a trending regime quite high when the price broke up past \$170 and therefore, by turning up the risk on trend-following techniques versus range-trading techniques or models, he should have benefited particularly well from the strong trend in copper.

The P&F charting technique is applicable for all traders or investors from the very short term to the very long term. If we wanted to analyse the potential trading regime on an intra-day basis we would just reduce the box size relative to the normal daily movement and if, as in the case of copper above, we wanted to analyse the very long term we would increase the box size to reflect this longer term analysis. Note, however, that still no concept of time is appearing in the analysis. It is a study of price only.

Figure 12.2 shows another example of using the Point and Figure method to identify both a range-trading regime and a trending regime. It is this “range leading to breakout” which is probably the most popular way to use point and figure analysis and is again the classical volatility breakout that is alluded to so often in this book. The chart shows the exchange rate between the British Pound and the Swiss Franc using a 0.005 box size and 3-box reversal. The price action in question takes place around late 2007 and early 2008 but the analyst does not need to know or be concerned with that. We can see that the market had been in a steep fall from around




Figure 12.2 GBPCHF – Point and Figure 0.005 × 3

an exchange rate of 2.30 to 2.13, whereupon a rally occurred taking the market back up to the previous zone from where a rally had occurred, around 2.1850. However, the market then failed to continue higher but instead fell back, crucially stopping and having a 3-box reversal around 2.1250. At this point, the Point and Figure/trading regime analyst should be alert to the possibility that a range is developing and, when the subsequent rally failed at the previous reversal zone of around 2.18, then the analyst should be quite confident that a range is in place.

One of the beauties of P&F charting is that it enables the analyst to identify support and resistance areas more clearly than using the modern bar, line, or candle charts. This is because by using a methodology that is only interested in when there has been a move or movement greater than a particular minimum (the box size), it eliminates the pinpoint highs and lows we are all accustomed to seeing on bar charts. In this sense it can be said to get rid of a lot of the noise that comes with things like stops and slippage in real life market trading, which might cause, for example, a previous high to be beaten marginally only for the market to then turn back down as the buy stops above the previous high turned out to be the only demand there. P&F charts

allow the analyst to consider support and resistance as a zone rather than a pinpoint number.

At 2.18 on the GBPCHF exchange rate, therefore, the actual high would have been different from the previous box reversal high but the P&F analysis means that the analyst can conclude that the reversal takes place “around” the same level. As with any analysis of a potential development of a range we have to be cognisant of course that we are dealing with the “potential” development of a range. At this point the analyst cannot possibly know how long the range will continue. He cannot predict it, but he can anticipate that the sideways trend in price will continue until evidence is presented that it has ended (until, as for Keynes, the facts change). He must assume that until evidence is presented to the contrary that the upper and lower bounds will remain in place and we can see from the P&F chart that the market presented the trader with at least six further opportunities to take advantage of the range-trading environment either by actively buying and selling GBPCHF or by factoring this range-trading information into his overall investment game plan, for example with regard to position sizing.

The longer this well-defined range was going on, of course, the more likely it was becoming that the eventual breakout of the range would provide an opportunity for an extended move. Sometimes the way I like to think about a range-trading environment, especially one that goes on for a while in relatively tight ranges, is that this is analogous to the market storing up energy and, once that energy is released, it should propel the market at a strong speed to the next level. Remember a range is equilibrium in price where supply and demand are in relative balance, so once that balance becomes upset then disequilibrium occurs.

At point B on the P&F chart we can see that the market started to slip down past the support zone of 2.1250 to 2.13 and, once it did so, it fell to 2.0550 without reversing 0.0150 (0.005 times 3). Thus the energy that had been stored up in the period of equilibrium trading was released with enough power to cause the market to fall relentlessly, in a trending regime for a period of time. In this example the box size we were using was 0.005 and given that the average daily range of GBPCHF tends to be around 0.02, then this analysis would have been suited only to shorter term traders or investors who wanted to identify short-term changes in trading regime. In actual fact the range in GBPCHF occurred over a period of around six weeks and the move down once that range was broken occurred in just four days. Therefore only the very short term focused would have taken advantage of this switch in trading regime from ranging to trending.

POINT AND FIGURE COUNTS

I have always been curious about so-called price objectives, meaning those levels where the price is expected to go to once a technical price pattern has been completed. As a student of technical market analysis I can testify to the fact that these

SUMMARY

In this chapter we have found how the classic Point and Figure chart can really help an analysis of trading regime. By only looking at the price movement and ignoring the time element, the analyst can focus much more intently on the most important element of trading regime (the price) and not be distracted by the passage of time when short-term pauses in a trend can sometimes be too quick in anticipating a trend exhaustion. Point and Figure charts allow us to neutralise a lot of the noise that exists in bar or candle charts because we are only interested in moves of a certain degree. We can eliminate a lot of the price spikes that occur along, for instance, horizontal range trading regimes where the bar chart analyst could have trouble identifying a proper price breakout. Finally, P&F charts can give us an idea of potential trend exhaustion zones by the use of horizontal counts.

price objectives generally hold true, but that they do so predominantly only as a first objective. What I mean is that the price objective is achieved but, generally speaking, the market does not necessarily just stop there and reverse back up. Many times it just keeps on going, leaving the person who closed his position at the so-called price objective feeling dejected. Take the example of a head and shoulders price pattern. The generally accepted wisdom is that the price “target” should be equal to the breakout level minus the distance of the “head” to the “neckline”. Generally, when these patterns occur this objective is achieved, but I cannot find much evidence that supports a view that the price just stops there. What we are talking about with these breakouts from ranges into trends is not an exact science and nor should it be. The idea and concept of price “targets” sounds nice and can provide, at the very least, a plan for *partially* exiting a position, but I have always been cautious about using price “targets” because, for me, the fact that the market is in disequilibrium and therefore in a trend means that the forces behind the market are quite strong. Why try to stand in front of that force and close a position that is running well for you? Agreed, a partial close might well be prudent but to fully close a position based on a price target seems to be a very bold strategy because you really run the risk of the market continuing to trend strongly in that direction.

Point and Figure charting can give the analyst a method of establishing a potential price target (at least in the first instance) when a market breaks out from a range-trading regime to a trending regime by utilising what is known as a “horizontal count”. The horizontal counting method is applied by the analyst counting the number of columns that occurred in the range and then applying that number to the breakout level in order to estimate the number of boxes that will probably occur (at least in the first instance) in the subsequent trending regime. If we take the example above of the GBPCHF exchange rate we can see from Figure 12.2 that there were 13 columns in the range-trading regime. Therefore, when the market broke down past the existing support of the range around 2.1250, the analyst could count the columns of the range (13) and apply that to the breakout level to obtain a price projection. In this example the box size is 0.005 and so 13 times 0.005 equals 0.065, meaning that the price projection would have been 2.1250 minus 0.065, which is 2.06. We can see from the chart that the market reversed, having hit a low box of 2.0550, so in this instance the initial price projection was fairly close to being spot on!

How do price projections help the trading regime analyst? I suppose they can help to determine a potential zone whereby a trending regime might exhaust itself and potentially turn into a range-trading regime and in this regard, when used in conjunction with other analysis techniques such as Bollinger Band Width, it might provide some use. My philosophical contention with price projections of course stems from the fact that I am a trend follower at heart, and so, for me, “predicting” where a trend might end is anathema. I prefer to listen to the markets and let them tell me where they are going so that I can anticipate trend changes rather than try to pick tops and bottoms of market moves.

13

Rate of Change and Divergence

Trading regime analysis is concerned with identifying periods when a market is trending or when it is range trading and thus helping the trader to enhance returns from those different volatility conditions or the change in them from one to the other. We have seen that identifying such periods can be done via various methods and indicators, but that the standard measure of historical volatility is not a brilliant indicator for such purposes. This is because historical volatility, which is defined as the standard deviation of the *changes* in market prices, might give a measure of the stability of prices but it does not really give a good indication of the dispersion (or range) of the actual prices themselves. However, what if we were to look at the pace of change of the prices themselves and not the standard deviation of the changes in prices? Would this indicator give the trading regime analyst a better steer on the current and likely future regime? Luckily, a standard indicator already exists for measuring this and it is commonly known as the rate of change indicator, or ROC.

The rate of change indicator is based on a simple calculation:

$$\text{ROC} = \frac{\text{Current period close} - \text{Close } N \text{ periods ago}}{\text{Close } N \text{ periods ago}} \times 100$$

Thus, the ROC is calculating the difference between a market's current price and the price N periods ago. It is a very similar indicator to a standard momentum indicator except that the ROC is expressed as a percentage, whereas the momentum indicator is normally expressed as a ratio.

As an example let's assume we want to calculate the 10-day ROC for a market that closed today at 23 and the close 10 days ago was 19. The ROC value today would be $((23 - 19)/19) * 100 = 21.05$. In this way, each day's ROC value takes different

inputs into the calculation and we can therefore form a time series chart such as shown in Figure 13.1.





Figure 13.1 EURUSD weekly – rate of change

Figure 13.1 shows a bar chart of the EURUSD exchange rate in the top chart and the 10-week ROC of that in the bottom chart. We can see that the ROC fluctuates between positive and negative numbers and that sometimes the absolute numbers are high and sometimes the absolute numbers are low. What is clear from the chart, however, is that when the ROC hits its extreme position either above or below the zero line, the market then stops going in its current position direction. This, of course, is a mathematical certainty in that the ROC is measuring the rate of change of the market and if the indicator is turning against its previous values it is because the market has turned against its previous values! However, the point to notice for the purposes of trading regime analysis is that when the ROC is at its extreme, it would appear to be a reasonable time to anticipate trend exhaustion and a strong possibility of a range-trading environment to develop.

The ROC can be used as a trading regime indicator in a very simple way by applying cut-off points to the indicator. It can be used as an indicator of when a market is entering into a trending period by simply looking at the level of the ROC as an indicator of trend strength. If the ROC moves over a predetermined level or

down through a predetermined level, then the analyst can be confident that there is a trend in action and, crucially, that it will *probably continue*. On the other side of the coin, if the indicator is staying inside the cut-offs then a range-trading regime is in operation.

Figure 13.2 shows the S&P 500 Index in the monthly time fractal and the 10-month ROC in the bottom chart. Now, let us suppose that you had decided that the parameters for deciding whether this market is in a trending mode is that if the ROC is above +10% or below -10%. The shaded areas on the chart represent when these conditions were met and we can see from the chart that the ROC was consistently above +10% from 1995 until 1998, when it had the briefest dip below 10%, and then remained above again until the start of 2000. Then, after the market went through a topping process the ROC dipped back below +10% before emerging on the other side below -10% where it stayed from around February 2001 to May 2003, with a brief bounce back above for a couple of months during that period. The indicator then signalled a trending regime from around June 2003 to June 2004 before the market went into a long period of crawling upwards in a very slow fashion until February 2007. This was an interesting period because, although it looks like a trending



Source: Bloomberg L.P.




Figure 13.2 S&P 500 monthly – rate of change

regime, on closer inspection it actually turned out to be a period of quite choppy price action when a lot of mechanical trend-following models underperformed, at least until late 2006.




Figure 13.3 Australian Dollar Index weekly – rate of change

Figure 13.3 shows the flow-weighted index of the Australian dollar exchange rate using the weekly time fractal from 2003 until 2007. This index is calculated in the same way as a trade-weighted exchange rate index but instead of the basket of currencies being weighted in accordance with the amount of physical trade (imports and exports) that the country does with other countries, the currencies are weighted by the amount of actual flow (or transactions) that is recorded in that currency *vis-à-vis* the index currency. So if Australia does a lot of physical trade with Indonesia, for example, but the Aussie dollar to Indonesian rupiah cross exchange rate is seldom directly traded, then it would not be given a big weight in this index. The flow-weighted index better reflects the actual trade that goes in the currency pairs themselves rather than the physical trade that goes on between countries and, in this regard, can be thought of as more of a realistic measure of a currency's valuation.

This data comes via Bloomberg from Citigroup, the institution that was responsible for inventing these useful indices.

We are analysing the Australian Dollar Index in the weekly time fractal and, therefore, we would need to choose a ROC cut-off in determining a trending or ranging regime that reflects such a fractal. The longer the time fractal the bigger the cut-off for deciding if the market is trending or ranging and the shorter the time fractal the smaller the cut-off should be. For the weekly time fractal in this index 5% seems like a logical number to use for such a cut-off and we can see from the chart that when the ROC was above +5% or below -5% then the market did in fact show some strong trending behaviour. It only went outside the 5% cut-offs a few times, and then only for a brief period, but when it did, the price action was a strong trend.

When the ROC stayed inside the cut-offs – meaning that the price was moving in a rate of change that was not over 5% positively or negatively – then, as we would expect, the market was staying in a range-trading regime. Again the reader might be thinking that this analysis is kind of circular. The market was not displaying a high ROC and therefore the market was trading in a range-trading fashion, and because it was trading in range-trading fashion it wasn't displaying a high ROC! The indicator is not even coincident with the price action but, given the calculation, it lags it. I can hear the quantitative analysis community say: What possible predictive value could an indicator such as this provide? Well, the thing about most traditional technical indicators is that they are a derivative of the price action, so they will all lag the price to some extent. Yes there are so-called neural networks and genetic algorithms that seek to predict price movement but, in my opinion, all they are doing, mostly, is taking price history and identifying cycles that have occurred in the past in order to anticipate potential future turning points. I am not saying they do not offer any value to the technical or trading regime analyst but, in my view, they are just asking a computer to do what a human brain can do for itself, given the same historical price information. Neural networks and genetic algorithms are said to “learn from the past” and be “continually learning” but isn't that exactly what a human price analyst does too?

Figure 13.4 shows the US Treasury 2-Year Yield in the monthly time fractal and the associated 10-month ROC in the lower chart. In this example we put a cut-off for determining whether the market is ranging or trending around 10% which seems like a logical number given the market (a yield) and the time fractal (monthly). We can see from the shaded areas where the ROC is above +10% or below -10%, and so these are the areas where the analyst would conclude that the market is in a trending mode. The areas that are not shaded represent the zones where the ROC is within the bands and where the analyst would conclude that the market is probably not trending but is ranging. We can see from the chart that the indicator with these cut-off levels did a good job in identifying the periods where a bond fund manager might be inclined to alter his risk profile or weightings between different investment styles accordingly.




Figure 13.4 US Treasury 2-Year Yield monthly – rate of change

SUMMARY

In this section we have looked at using the simple rate of change of a market price as a trading regime or volatility indicator. We have seen that by thinking logically about tolerance limits we can identify when the rate of change is considered to be slow with the market in a range-trading regime, and also when the rate of change is considered to be fast with the market being in a trending regime. This simple analysis, although not predictive, is able on many occasions to anticipate persistence in the regime and in this way it allows the analyst or trader to factor the regime analysis into his work.

DIVERGENCE – ANOTHER MEASUREMENT OF TREND STRENGTH

My motivation for studying the philosophy of trading regime analysis was borne out of the realisation that trend following works brilliantly some of the time, but works badly at other times. Markets trend and then range. I wanted to know not only why this happened, but how I could anticipate the periods of trending and range trading.

In this regard the question really boils down to one of the strength of a trend. If a trend is strong then it should be anticipated that it will continue for the next period, but if a trend is weak then equally it should be anticipated that the “trendiness” of that trend is breaking down and a period of choppiness or range trading is quite likely.

There are various methods of analysing the strength of a trend and one of the most popular is the ADX examined in Chapter 11. However, another commonly used approach, albeit for a slightly different purpose, is to look for what is called divergences between the price of the market and a measure of momentum of that market.

Momentum measures can be called different things (rate of change, momentum, relative strength, stochastics, moving average convergence divergence, to name a few) but they all have one thing in common – they should all be strong when the underlying market is strong and they should all be weak when the underlying market is weak. If the market shows a situation whereby the price itself has, for instance, continued to strengthen whereas the momentum indicator has not continued to strengthen, then the underlying strength in the market up move has undoubtedly weakened. By the same token, if the market price continued to weaken and yet the momentum indicator has not continued to show weakness, then the underlying strength in the down move has undoubtedly weakened. The momentum indicator shows the power and energy behind the market, and if it is not marching to the same tune as the market price then something is wrong with the market. The analogy that was used with regard to the rising wedge in Chapter 4 is extremely appropriate with regard to divergence. If someone is walking up a down escalator, they are constantly fighting against gravity, and at some stage their energy levels will fall to the point where that gravity will take them down. This is how divergence between the price and momentum can act as an eventual trend reversal, and this aspect of trend reversal is probably the way that divergence is most widely used. However, perhaps what is of more interest to the trading regime analyst is the *process* of divergence itself and how that can lead to a period of choppy range trading before, and sometimes in place of, a trend reversal.

One of the most crucial aspects of understanding the markets is that trend exhaustion does not have to mean trend reversal. A trend can exhaust itself and go sideways for some time before taking off again in the same direction. In fact this happens more often than not. A lot of the time we look at trend exhaustion methods and anticipate a quick reversal in the previous trend when, perhaps, we would be better placed to anticipate a sideways to slight reversal movement. This is one reason why taking advantage of trend exhaustion can be better done via option call or put spreads rather than via the cash or futures market. Yes, a sharp trend reversal can occur after distinct trend exhaustion measurements but experience would suggest that it is best not to bank on that outcome only. In this regard, bullish or bearish divergence in relation to a momentum indicator can go on for some time and the analyst or trader who is expecting the initial divergence to result in a trend reversal will, on many occasions, be disappointed as the market continues to show divergence. However, if the analyst

moved up much and could only register a reading of 65.49. So we have a situation where the market has made a new high but the momentum indicator has not, and this would indicate that the underlying strength in the Canadian dollar is not as strong as it was back in the summer of 2005. This is what is known as classic bearish divergence but, as mentioned, the fact that it is called bearish divergence does not have to mean that the market will fall. The presence of divergence, in this case bearish, only means that the *previous trend has weakened* – it does not have to mean that it will reverse. So is it therefore not just as likely that the up trend rather than reversing into another down trend will instead dissipate into a range trading environment?

We can see from the Canadian Dollar Index chart that over the next few weeks and, indeed, months the price itself continued to make higher highs but the RSI was not able to rise past the original high reading in September 2005. Moreover, the trend itself became noticeably more choppy and volatile in the sense of the rate of change of prices rather than the measurement of price dispersion. Whereas the market had moved up from 119 to 131 in the space of four months in a strong trending regime, the market then spent its time oscillating in sideways to up ranges for the next year. During this period most intermediate term trend-following systems would have underperformed badly in the Canadian dollar, whereas if the trader or investment manager had pursued a more mean-reverting strategy he would have outperformed.

I used to think that technical analysis was possibly not applicable to certain markets, such as very illiquid stocks that hardly ever trade, but I also thought it was not applicable to certain markets that were in some sense controlled or manipulated by a controlling body. Oil might be one example, but the power of OPEC (Organisation of the Petroleum Exporting Countries) in controlling the free market in oil is somewhat questionable. Another market I thought, in my younger days, might be less applicable for technical market analysis is the short-term interest rate market because I thought it was the central banks that essentially controlled and manipulated this market via their verbal intervention or guidance, or by their direct actions. Over the years I have come to accept that this is actually not the case and that a technical market analysis of the short interest rate futures is just as valid as any other market. In some respects it could even be more valid because the sheer volume of the market in, for instance, the 3-month US interest rate future dwarfs most other markets. There is a very large and diverse crowd that trades the Eurodollars, so an analysis that examines that crowd psychology is most valid indeed. In fact, in recent years it has become quite evident that the Federal Reserve (and other central banks) actually moves interest rates to reflect what the market is anticipating rather than shock the market with what it is not. In that sense the market could be said to lead the Federal Reserve's decision but, to be fair to our fundamentalist friends, it could also be said that the market might not move to price in its expectations in the first place without the verbal guidance of the central bank. It is a sort of circular argument and not one, trust me, you want to get involved with if you want to keep your sanity! Nevertheless, there is a massive, free market in short-term interest rate futures, and so looking at

looked a little more closely at the price action he would see that although the previous move has been extended, the underlying regime of that move has subtly changed from trending to range trading. Let's look at an example to examine this in more detail.

Figure 13.5 shows the bar chart of the Canadian dollar flow weighted index (courtesy of CitiFX) using the weekly time fractal from December 2004 until the end of 2006 and the 14-week relative strength index in the lower chart. The relative strength index (RSI) is a very popular momentum oscillator, which measures the velocity of directional price movements. Looking at the chart we can see that in the summer of 2005 the Canadian Dollar Index was making new highs in a solid uptrend from around 120 in May before topping out at 131 at the end of September. On that week the RSI gave a reading of 72.01, indicating that the market, at that time, could have been considered overbought. Being overbought however does not necessarily mean that the market has to fall and instead it could go sideways for a while in order to work off that overbought condition. The market did, in fact, go sideways over the next few weeks before resuming the up move making a new high above the previous high at 133 on the week of the 16 December. However, at this time the RSI had not




Figure 13.5 USDCAD weekly and RSI

the mass psychology of such a market should be helpful for the analyst who wants to anticipate the future trading regime.

Figure 13.6 shows the September 2008 Eurodollar futures contract using the weekly time fractal from January 2007 until April 2008 and the associated (9, 5, 5, 3 parameter) slow stochastic oscillator in the lower chart. We can see from the chart that at the end of June 2007 the market entered into a strong trend that also led the slow stochastic higher into September and caused the momentum indicator to make a higher high coincident with the price at the end of November 2007. However, when the market subsequently rallied in price into January and early February 2008, the slow stochastic could not reach the highs of November, and this was a flashing warning to the analyst that bearish divergence was occurring with the Eurodollar market *losing its upward strength rapidly*. In fact, the price action on the week of 25 January was a rare “long legged doji” in Japanese candlestick terms (a doji that has a very wide range signalling extreme uncertainty) and so that in itself was a big warning to the analyst that the uptrend was potentially at exhaustion point.




Figure 13.6 Eurodollar Sep 08 weekly and slow stochastics

By early February, therefore, the trading regime analyst would be fairly confident that, due to the market losing its strength, a period of range trading was quite a high probability. We can see from the chart that the market subsequently entered a period

of sideways trading with opportunities to profit from this range by selling upside volatility or by raising the risk on mean-reverting models.

SUMMARY

This section has taken a brief look at the periods when a market trend starts to lose its strength as defined by divergence between the market price and an oscillating or momentum indicator. We have seen that identifying such periods can help the trading regime analyst to identify periods of trend exhaustion and potential range-trading regimes. In the next chapter we shall look at an indicator that is similar to these momentum indicators but can be used in a slightly different manner in order to anticipate trading regime.

14

Williams %R

“It’s not whether you win or lose – but whether I win or lose.”

Sandy Lyle

Trading or investing is a lot like sport. There is an arbiter or a scoreboard in the form of your performance or P&L, but unlike the majority of arbiters in sports this one is completely and utterly objective! You can’t turn round to the referee and dispute a call or a judgement in trading. The P&L is what it is and there is no subjectivity involved. As I have alluded to previously, this makes trading very similar to golf in that the trader or investor, just like the golfer, is forced to take sole responsibility for the outcome of the trade or shot, be it good or bad. There is no place to hide from this responsibility. Also, just like sports, investing is competitive with league tables of the best and worst performing funds.

It is little wonder why many sports-minded individuals are attracted to the financial markets trading or investing realm and, notwithstanding the fact that fund managers are of course competing against each other for the best performance, sooner or later someone was going to come up with actual trading competitions using real money. A lot of these trading or investing competitions are used in schools and colleges, where the entrants get a great insight into the cut and thrust of the markets. However most of these use “paper” or fictitious money and, while that can give you a very good understanding of what it is like to trade or invest for real, it is not until you are managing real money, especially your own, that you experience the full gamut of raw emotions that come with the territory and which are the hardest aspects of our “game” to come to terms with. There are a number of trading competitions using real money but the jewel in the crown is recognised to be the World Cup of Trading run by Robbins Trading in Chicago. Over the years there have been many winners of this prestigious event (a competition of return over a one-year horizon) that have become very well known in the industry and beyond, but one figure stands out because of the stupendous returns that won him the competition.

Larry Williams is a trading legend. In 1987 he entered the World Cup Trading Championship and won with a one-year return of over 11,000%. Eleven thousand percent! His starting futures account of \$10,000 was turned into one of over \$1.1 million, but perhaps the most interesting aspect about this was the fact that, before the stock market crash in October 1987, the account was actually over \$2m. Williams's return that year has never been close to being beaten in the competition. In fact, one of the few people to have had multiple returns like this was Williams's daughter in 1997 with a return of 1,000%. These returns are real and audited.

An objective analyst might think that the Williams family have something to offer the trading world and, of course, Mr Williams has authored some great books on the subject of trading and investing, making his techniques and methods well known. His most famous indicator is probably the Williams %R that can be found in most technical analysis software packages, and that is what I want to focus on in this section.

The Williams %R calculates the difference between a market's most recent closing price to its highest high price relative to its price range over a given time period. In this regard the Williams %R is similar to stochastics, but rather than use the lowest low as the reference from which the current close is compared as stochastics do, it uses the highest high for the reference. Like stochastics, however, the Williams %R indicator is typically used to identify overbought and oversold conditions in the market and, in this regard, it is *generally* considered to be a contrarian indicator. However, I would like to focus on using the Williams %R as a volatility breakout indicator rather than an indicator that helps the trader to take advantage of range-trading regimes.

The Williams %R is calculated as follows:

$$\text{Williams \%R} = \frac{(\text{Highest high in } N \text{ periods} - \text{Most recent close}) \times -100}{\text{Highest high in } N \text{ periods} - \text{Lowest low in } N \text{ periods}}$$

The Williams %R is therefore very similar to the %K in stochastics, which is calculated using the lowest low rather than the highest high. It is a relative strength measure of where the closing price is relative to the highest high and the range over the period. It will therefore, like stochastics, measure where conditions are overbought or oversold in relation to recent ranges. However, this is a fast indicator and does not spend too much time within the body of the values unless there is a well-defined range. Instead it tends to go very quickly to the extreme values and, when there is continued movement in that direction, it stays there. This aspect of an oscillator staying overbought or oversold as the market trends in one direction or another is well known, but the Williams %R gets into these overbought or oversold territories more quickly than the traditional oscillator indicators and it is therefore quite useful in pre-empting range or volatility breakouts.

Figure 14.1 shows the Williams %R indicator in action using the monthly time fractal on the OEX Index in the USA from the early 1980s until the late 1990s. The

In December 1984 the OEX Index closed at 82.54 and then entered into a multi-month uptrend (when the indicator did not move back down through the -30 level) culminating in the 1987 stock market crash in October 1987. In August 1987 the OEX had reached a high of 166.14, a 101% return from the level that was indicated as the start of a potential uptrend.

Even though the Williams %R indicator is a fast-moving indicator, it was not quick enough on this monthly time fractal to anticipate the 1987 crash and, of course, had you held a long position based on this indicator from 1984, the 1987 crash would have eroded the return to something like 30%. However, in my opinion, market analysis and trading regime analysis in particular should be a *multi-fractal* and semi-discretionary exercise. What I mean is that this indicator, as with most other trading regime indicators, should be exactly what it says it is – an indicator – an indicator of a potential trend. How the trader, fund manager, analyst wants to take advantage of that trend should be a separate decision based, in my opinion, on one time fractal down from the indicator. What I mean is that rather than using the Williams %R as a stand-alone trading system I would prefer to use it as part of a process whereby it identifies the potential trends and I then use something like moving averages one time fractal down in order to take advantage of those trends. So in this case, when we are looking at the monthly time fractal for the Williams %R turning higher in 1984, we would think of taking advantage of this potential uptrend using, perhaps, weekly time fractal moving averages or even a weekly fractal of the Williams %R itself. Doing this means that we might not be “in” the uptrend over its entire life but we will probably be “out” before something like the 1987 crash happens.

Figure 14.2 shows the weekly time fractal Williams %R on the OEX over the time period in question and we can see that the weekly indicator, having been indicating a bullish trend from early June when the index was at 144.11, actually turned back into the neutral zone in early September 1987 when the index was trading at 156.46. So had the *actual market positions* been taken using this multi-time fractal methodology, then being long the OEX during the 1987 crash would have been avoided.

The Williams %R is an extremely fast indicator and it can actually go from indicating one extreme in one time bar (bar chart of any time fractal) to the other extreme in the next time bar. This is both good and bad. It is good because the indicator should give the trader a lead on the anticipated next move in the market but it is bad because acceptance of dramatic change like this can sometimes be difficult for our human brains. Humans do not generally feel comfortable with *quick* change but rather prefer change to be slower and more evolved, and this is in fact why using objective measures of change like technical analysis indicators help to *force the process of change* on human brains when the natural inclination is to resist. The sharp oscillations in the Williams %R indicator are therefore hard to accept but, as we can see from the chart, these quick changes can give a good lead in anticipating changes in volatility conditions.

Overall, then, we can see that the Williams %R can give the analyst a reasonably good indication of when the market is about to experience a volatility breakout and




Figure 14.1 OEX monthly and Williams %R

OEX is a good barometer of the general health of the US economy as it consists of 100 of the largest US companies upon which options are tradable. The Williams %R indicator, as mentioned, uses overbought and oversold levels generally considered to be at -30 and -70 . However, for the purposes of the trading regime analyst we want to use the speed of this indicator to identify when it goes into “overbought” territory to identify a potential uptrend, when it goes into “oversold” territory to identify a potential downtrend, and when it is staying within these boundaries to identify a potential range-trading period.

We can see from the chart that, having been indicating a downtrend and then a range, in December 1984 the indicator moved above the -30 level into overbought territory. Remember, this is a fast indicator and we are turning the normal oscillator logic on its head by using it to give us a lead indicator of trending moves, so when the indicator moves up above -30 the trading regime analyst should be thinking that the market is probably entering an uptrend regime. What the indicator is telling us is that volatility is ticking up (as measured by the most recent close to the highest high of a given look back period) and that, in all probability, this will lead to a volatility breakout and a trend in the price.




Figure 14.2 OEX weekly and Williams %R

a subsequent trend. However, like all indicators it is not infallible and has to be used in conjunction with other analyses in order to get the full benefit.

SUMMARY

Much of the trading regime analysis we have looked at thus far has a basis in volatility breakout or trend-following type of analysis and, in the next chapter, we will look at an example of probably the most famous of all trend-following techniques and how this encompasses an analysis of trading regime.

15

Donchian Channels

“Get those brokers back in here. Turn those machines back on!”

In the 1983 movie *Trading Places*, with Dan Aykroyd and Eddie Murphy, the plot revolves around a bet that two old-time market players have with each other over whether a highly respectable futures market executive (Aykroyd) could swap places with a down on his luck vagrant beggar (Murphy) and not turn to a life of crime. The comedy story then goes on to climax with the two opposite characters getting together to take their revenge on the old timers who had their fun with the bet. Legend has it that the movie inspired (or was inspired by, depending on who you speak to!) a real life bet, or challenge, that took place in the futures markets between two market traders around the same time; and the challenge was over the issue of whether successful market trading was an inherent skill or whether it could be learned. The two people in question were Richard Dennis and his business partner William Eckhardt who ran a successful trading operation and, after Dennis had visited a turtle farm in Singapore, he was said to have exclaimed that they could “breed traders just like they breed turtles”. It would appear that there was a disagreement over this, which led to the challenge, and the legend of the “Turtles” was born.

An ad-wanted was put in the newspaper looking for prospective traders and they eventually chose people with diverse experience about the markets, whereupon they put them into their training programme. They were given a strict set of rules to follow regarding trade entry, trade exit as well as money management and they were let loose on the markets with capital backing from Dennis’ company. The group became known as the Turtles and most of them went on to be very successful traders.

The story of the Turtles is well known in the trading and investment world and their fame gave a big boost to their trading method being introduced and accepted outside the futures markets. The method that Dennis taught them was trend following, which

was until then mainly used in the futures markets but has gained wider recognition in the cash markets since, and especially since Michael Covel of www.turtletrader.com has written a very good book on trend following. For the purposes of this book, however, we want to concentrate on the actual strategy that people like the Turtles have applied, and see why that helps us with our mission of trading regime analysis.

There are many different brands of trend-following systems but they can all be split into one of two categories. They can be a “continuous” system whereby the system always has a position in the market either long or short. For example, a trading system that is to open a long position when the price crosses above a moving average, and to close the long and open a short position when the price crosses below the moving average, would be an example of a continuous trend-following system. Or they can be “discontinuous”, whereby the system can go through periods of having no position in the market. For example, a system that opens a long position if the market price is 2% or more above a moving average and opens a short position if the market price is 2% or more below a moving average can go through periods when there is no position in the market. Moving average envelopes can easily be made into a systematic trend-following system as described.

Richard Donchian is the person regarded as the founder of trend following as a modern investment philosophy and he was a pioneer in the field of managed futures, the area of the fund management industry that manages retail investments in funds where the primary investments are made in the futures markets. He did more than anyone to develop a systematic approach to futures trading which he allied to promoting quite a conservative approach to managing futures funds.

Donchian was born in 1905 in America and became interested in technical analysis after being made aware of Jesse Livermore, the great tape reader and focus of the book *Reminiscences of a Stock Operator*. During the 1930s he worked in Wall Street before becoming involved in World War II when he participated in the invasion of Sicily, and later served as an Air Force Statistical Control Officer in the Pentagon. After the war he developed more of an interest in actively trading rather than analysing the market and developed a rules-based technical system. He established the trading firm Futures Inc that was then one of the first publicly held commodity funds.

Donchian's rules-based system was developed on his observations that commodity markets tend to move in long, sustained moves and it therefore became known as trend following. His work is the basis of most trend-following systems and, although his studies and work were very broad in nature, his “trademark” for trend followers came up with the idea of trend “channels”, now widely referred to as Donchian Channels.

Donchian produced a set of rules and guides (first published in 1934 and which can be found on the internet), the most important of which he believed was number three on his General Guides: “*LIMIT LOSSES, ride profits – irrespective of all other rules.*” This really sums up the overriding ethos of Donchian's work in that, to him,

the trend was, in the final analysis, the most basic and important part of any market analysis.

Donchian's work was the basis of the Turtle rules and more specifically it was the so-called "weekly rule" that became the anchor of the methodology. The test results for this system were published as early as 1970, and it was found to be the most profitable system then known. The weekly rule, in its simplest form, buys when prices reach a new four-week high and sells when prices reach a new four-week low. A new four-week high means that prices have exceeded the highest level they have reached over the past four weeks. Likewise, a four-week new low means prices are trading lower than they have at any time over the past four weeks. This system is always in the market and so it is a continuous trend-following system. Known simply as the four-week rule, this is the system initially designed and used by Donchian and it was only later modified to take account of the concept of discontinuous trend following whereby the system would have no commitments in the market during certain times.

Discontinuous trend-following systems are of significant interest to the trading regime analyst because they are actually intending to take account of periods when there is *sideways* movement (or choppy price action) in the market by not making a trading decision. Essentially they are using a methodology to identify a range-trading regime and therefore not have the trend-following system take any positions during those periods. In this regard a discontinuous trend-following system can be viewed as a complete trading system because it at least tries to take account of the probability that there will be times when the system methodology will not work.

The detailed rules that the Turtle group applied are relatively well known (again, you can access them on the internet) but it is the logic behind these detailed rules that interests us as trading regime analysts, and therefore the rules presented here, although based on Donchian Channels and the original Turtle method, are by no means the full Turtle system in detail. The logic, however, is the same and this is why I am referring to them as a "Turtle Type System". Turtle aficionados will no doubt castigate me for taking the name in vain, but the point I want to make is that the philosophy of the Turtle system was to identify when a trending regime was happening and when a non-trending or range-trading regime was happening.

Rules for trade entry:

- Open long if the current closing price is the highest closing price for the last 20 periods.
- Open short if the current closing price is the lowest closing price for the past 20 periods.

Rules for trade exit:

- Close long if the current closing price is the lowest closing price for the last 10 periods.
- Close short if the current closing price is the highest closing price for the last 10 periods.

Here we have two rules for entering and two rules for exiting market positions and, in the absence of any other information or rules, this is a basic discontinuous trend-following system. What it means is that the system will only ever open a long position if the market is breaking out above the closing high of the last 20 periods and will only ever open a short position if the market is breaking down below the closing low of the last 20 periods. When the system is in a long position it will exit that long position if the closing price is the lowest of the last 10 periods and vice versa for a short position. So we can see that there will be times when the system will not have an open position in that, for example, the system could have closed long or short positions but, because the market has not made any fresh 20-period closing highs or lows, it will not have initiated any fresh long or short positions. The system is designed this way so that it minimises the times when false signals or “whipsaw” occurs in range bound or “choppy” market conditions. Only when there is a fresh closing 20-period high or low will the system initiate a trend-following position. Therefore, the inference from the system logic is that if the market is not making fresh 20-period highs or lows then the market is, to all intents and purposes, in a range bound market. The system is programmed to define a trending regime (fresh 20-period highs or lows) and a range-trading regime (no fresh 20-period highs or lows).

So we should be able to take a time series of a discontinuous trend-following system like this and see the periods at which the system was identifying a trending regime and the periods at which the system was identifying a range-trading regime. Figure 15.1 shows just such a time series chart, and the periods of trending regime and ranging regime identified by this simple system.




Figure 15.1 USDCHF weekly – Turtle type system

The original Turtle methodology was only based on daily data to establish the system rules. For the longer term trading regime analyst, however, using the daily time fractal is not long enough to give a good analysis of the potential longer term trading

regime and so it is wise to apply the same methodology to weekly and monthly time fractals to give a longer term analysis.

Figure 15.1 shows the bar chart of the US dollar exchange rate with the Swiss franc in the weekly time fractal. The black lines represent the highest close over the last 20 bars and the lowest close over the last 20 bars. This is the basis of a Turtle type trend-following system or a discontinuous trend-following system (but note that the original Turtle rules referred to the daily time fractal only). The logic is that if the market is making new 20-bar highs then it is in an uptrend and if it is making new 20-bar lows then it is in a downtrend. If the market is not making new 20-bar highs or lows then it is deemed to be in a neutral state or a range-trading regime.

The shaded areas in the chart show those periods when the market was not making new 20-bar highs or lows and was therefore, according to the logic of the system, in a range-trading regime. Only when the lower black line is going down or when the upper black line is going up is there deemed to be a trend. If both the bands are narrowing in, then it is signalling that the price action is trading in an ever-decreasing range. We can see that, for this data set, the market was in a range-trading regime by far the majority of the time. When the market did start to make new 20-bar highs or lows it did so only very briefly before returning to a ranging environment.

If we look at the same exchange rate, but this time in the expanded monthly time fractal, we can see that the same methodology can be appropriate too. Figure 15.2 shows the USDCHF monthly chart and we can see that there were again quite a lot of periods when the system was signalling a range-trading regime and, when it did, the market did trade in quite a well-defined range. This of course is simply a reflection of the system parameters but the point is that the trending or range-trading regimes can exist long enough for value to be extracted by a trader or investment




Figure 15.2 USDCHF monthly – Turtle type system

manager, especially if he is using the higher time fractal for volatility analysis and a lower time fractal for actual position commitments. In the period from around December 2005 to August 2007 the market traded in a range of between 1.31 and 1.19, giving the trading regime analyst or trader many opportunities to add value from adjusting trading strategy or factoring that information into their overall money management plan.

However, in the monthly time fractal on this data set there were times when the methodology signalled a strong trending regime that lasted for a number of months. In the period from around May 2002 to June 2003 the market was making fresh 20-month lows and therefore the system was signalling that the market was experiencing a strong downtrend. In that time period the market dropped from around 1.55 to around 1.30, giving good opportunities to benefit from this downtrend over all time fractals.

The Turtle rules incorporated a two-speed trend-following system in that the rules were based around new highs or lows on a short time period and new highs or lows on a longer time period. They used 20-day highs and lows to determine the shorter term trend, but they also used 55-day highs and lows to determine the longer term trend. In this way, the system was setting up to capture long-term and intermediate term trends. Note that there is nothing magic about the numbers the Turtles used. Like moving averages or any other trend-following parameter, it's really not about *optimising* a system to create the best performing system over a historic data set; it is much more about defining a sensible logic and having the *discipline* to stick with that logic through good times and bad. The specific entry and exit points of a trend-following system are not nearly as important as the *logic of running winners and cutting losers* to create a long-term mathematical expectation of the average gain being higher than the average loss.

THE MOST IMPORTANT PART OF ANY TRADING SYSTEM

The entry and exit rules were only a part of the original Turtle trend-following system. Another major part of the system, and arguably the most important element of this and any other system, were the money management rules. Now, thanks to the fact that there is a dearth of literature on the subject many people think of “money management” as stop-loss and take profit levels. In actual fact, as was mentioned in Chapter 1, the most important element of what is known as money management is the question of what *size of risk* to commit to a trading decision because this is the variable that will have the biggest determinant on your overall profit or loss. The position or bet size will have the most profound effect on any trading or investment performance yet it seems to be the subject that is least talked about or analysed in the markets today. We all focus on the “*when to get in*” and “*when to get out*” questions but, in reality, the amount of risk we have will continuously relegate these question to second place. Even if we time the market to perfection, if we have only a small amount of risk on the trade then it really doesn't matter that much, and if we didn't

time the market at all well and we had a lot of risk on board, it can become disastrous. As many people who have traded or invested in the markets will know, one of the psychological rules of the game is that “when it is going against you, you always have too much risk on, and when it is going your way, you always have too little risk on!” Position sizing is quite simply the biggest determinant of whether you will be profitable or not as a trader or investor over the long term.

The Turtle method had very specific rules on position sizing, and if you want to see the details then simply look it up on the internet, but yet again it is the logic and concept behind these rules that we are interested in when it comes to trading regime analysis. The position sizing rules for the Turtles was based on a measure of volatility called the Average True Range with the simple logic that when there was low volatility (as measured by the Average True Range) the position size would increase, and when there was high volatility (as measured by the Average True Range) the position size would decrease. So, given that the system itself can be thought of as a volatility breakout type of model, the biggest positions are going to come when the market is coiled up and ready for a volatility breakout and the smallest positions are going to come when the market has had an extended move and the risk of choppy or range bound market conditions have increased.

THE AVERAGE TRUE RANGE

The Average True Range (ATR) is a measure of volatility introduced by J. Welles Wilder in his book *New Concepts in Technical Trading Systems*. It is based on a calculation of what Wilder called the True Range of the market, which is defined in Chapter 11 of this book on the ADX. The ATR is a moving average of the True Range and although Wilder suggested a default period of 14, any number can be used depending on how sensitive or insensitive the analyst wants the indicator to be. Wilder originally developed the ATR for commodities but the indicator can also be used for stocks and other markets.

The True Range is calculating the range of market prices in any time period be that hourly, daily, monthly and so on. It is therefore measuring how narrow or how expansive the price range in that particular period has been and so the ATR simply measures how narrow or how expansive price ranges have been over a certain number of periods. It is therefore a very good indicator of how price ranges are expanding and contracting on a per period basis. This is different from a methodology that looks at the standard deviation of prices, such as Bollinger Band Width, because the underlying calculation in the ATR is based on the observations of price ranges between highs and lows, whereas the standard deviation calculations are based on the closing prices of each period. Nevertheless, the ATR is still a very good indicator of market volatility because the methodology is consistent with what one would expect of market behaviour during periods of range trading and periods of trends.

When markets are trading in a range we would expect the gap between highs and lows in each trading period (hourly, daily, etc) to get narrower as market participants become less and less interested and therefore likely to commit less to the market. This dampens volatility as measured by the ATR. On the other hand, when markets are trading in a trend we would expect the gap between highs and lows in each trading period to get bigger as market participants become more and more interested and therefore likely to commit more to the market. This increases volatility, as measured by the ATR.

The logic of using the ATR as a trading regime indicator is similar to others. After range contraction comes range expansion, and after range expansion comes range contraction. That is, when the high/low ranges of trading periods have been extremely small, it is more likely than not that the ranges will increase resulting in a trending regime, and when the high/low ranges of trading periods have been extremely large, it is more likely than not that the ranges will decrease, resulting in a range-trading regime.

Wilder's starting point, the True Range, was an idea he came up with essentially because he was mainly involved in commodities trading in the 1970s when he published his book. Commodities like gold, soybeans and cotton can be very volatile in the sense that they can experience so-called "limit moves", meaning that if a commodity moves up or down by a specified amount, then trading is suspended for that day and the market will not trade until the next day. Sometimes, the market will open the next day limit up or limit down, and so no trading will be done at all. The price recording for that day will be one single price (i.e. open, low, high and close are all the same). There can also be a lot of gaps in commodities price data, meaning that a market in an uptrend will close at, for example, 10 and the next day gap up at the opening to 11 before continuing with the uptrend. Both of these problems meant that calculating the range of prices for each particular day was problematic in that just looking at the high to low range for the particular trading session would not be an accurate reflection of the *true* range if there had been, for example, a gap from the previous close to the current open. This was the motivation for Wilder to consider the concept of the True Range.

One of the problems with the Average True Range is that over the course of time the normal ranges that a market will experience can change structurally higher or lower. What constitutes a "normal" range could be influenced over time by the amount of interest or participants in the market and structural issues such as government policy and regulation. In a sense this is a problem with any volatility indicator because volatility can become structurally higher or lower over the course of time, and this is why it is useful to normalise the data on very long-term volatility charts with the aid of averages or other techniques. However, some indicators, like those based on the standard deviation of prices, can give a more consistent indication of when volatility is at a low or a high level. Nevertheless, the prognosis of what constitutes a high or low level of volatility remains quite a subjective process and we

illiquid and small markets like some emerging market or small cap stocks, then we know that quiet, apathetic and inactive markets are very unlikely to persist forever and so, by early 2002 with the ATR being at historic lows, the trading regime analyst would have been on the lookout for that volatility to change and for a regime switch from range trading to trending to probably occur.

We can see from the chart that in February 2002 the Norwegian krone started to move higher out of the range-trading environment, and that in April 2002 the ATR started to respond by moving higher again. In fact the krone moved higher in a strong trending move over the next few months and, until August 2002, the ATR was also getting higher and higher. In April 2002 the ATR peaked out after the underlying market had hit a short-term top, but the market continued to move higher thereafter albeit in a choppy fashion. So in this example, although the extremely low level of the ATR was a good lead indicator of the trending regime that was to follow, it is not that clear that a high to falling ATR led to a range-trading regime at least initially.

One issue with the ATR as a trading regime indicator is that sometimes a persistent trend can develop that has, at least in the early to mid-stages of the trend, stable to even falling price ranges on each bar and that ranges can develop that have rising bar price ranges. While this goes against the general logic of quiet (and therefore ranging) markets having small ranges and noisy (and therefore trending) markets having higher ranges, it is still consistent with our logic at the extremes. Generally, we expect volatility to be extremely high towards the end of trends and we expect volatility to be extremely low at the start of trends. What volatility does during the course of the trends or ranges, as measured by the average bar price ranges, will be different from what volatility does during the course of ranges as measured by the standard deviation-based indicators. In those indicators it is the range of price expansion and contraction over the last n periods that is being measured, whereas with the ATR it is the average bar price range that is being measured. In this respect the ATR acts more like a measure of historical volatility where the standard deviation of the *change* in prices is being recorded. However, the ATR still acts as a reasonable measure of “chop” in the market. What I mean by this is that when the ATR is at a high level, the risks increase that the market will act in a very choppy fashion, and when the ATR is at a low level, the odds on the market chopping around will be quite low. From a risk management perspective it is important to be able to quantify your risk and the ATR can give the analyst or trader an insight into how likely it is the market will fluctuate on a period-by-period basis. Many people use the ATR as a money management tool in terms of calculating stop-loss levels, and this idea was also incorporated into the original Turtle trading rules.

Figure 15.4 shows the bar chart of the OEX Index in the USA using the weekly time fractal with the 12-week ATR in the bottom chart. The chart captures the price action from 1998 to 2003. We can see from the chart that periods when the ATR has risen and is peaking out coincide with choppy market conditions or general range trading such as in early 2000 and again from August 2002. It is also clear that periods




Figure 15.3 NOK Index and ATR weekly

really have to include the direction of volatility to enable us to add that to the overall trading regime analysis.

Figure 15.3 shows the bar chart of the flow-weighted index of the Norwegian krone using the weekly time fractal from 2000 to 2004. The bottom chart shows the 12-week ATR, meaning that it is the average of each week's true range over the last 12 weeks. We can see that from 2000 the Norwegian krone was trading in a distinct range-trading regime and the ATR was generally stable to falling. By early 2002 the ATR had fallen below the 1.00 level, which had marked the previous low levels as the recent weekly ranges of the Norwegian Krone Index had become ever smaller. The psychology of the market had become one of apathy, with interest and activity being very low. We can conclude this because of the falling ranges. If interest and activity were high, then the ranges would have been larger as more volume would have led to greater moves in the market. Unless we are talking about extremely




Figure 15.4 OEX Index and ATR weekly

when the ATR has fallen and is recovering coincide with periods of trending regimes such as from September 2000 and April 2002. In April 2002, in particular, the ATR had fallen to a very low level and this period of “calm” in the market preceded the “storm” that followed when the index lost over 25% of its value in just three months.

SUMMARY

The Average True Range is a good indicator of volatility in the market in terms of the general level of fear/greed or complacency in that high readings indicate potential trend-ending points and low readings indicate potential trend-starting points. However, it does appear to be more inconsistent than other measures when applied to identification of trading regime and, in my opinion, it should therefore only be used in conjunction with other indicators as part of an overall strategy.

STARC BANDS

Like other measures of volatility, the ATR is a very portable indicator and it was used by Manning Stoller in the 1980s to create what he called Stoller Average Range Channels or STARC Bands. The logic of STARC Bands is essentially the same as with Bollinger Bands, but instead of using the standard deviation to calculate the width of the bands they use the Average True Range. So the bands will expand and contract in line with the expansion and contraction of the ATR.

Plotting two bands around a short-term simple moving average of an underlying asset's price creates STARC Bands. Adding a value of the ATR to the moving average creates the upper band while subtracting a value of the ATR from the moving average creates the lower band.

Upper STARC band = Simple moving average + ATR

Lower STARC band = Simple moving average – ATR

The ATR is generally multiplied by a user-specific multiplier factor before being added and subtracted from the moving average, and although the original default of the average was a six-period average the user can of course experiment, as with all indicators, with what periods have the right logic and feel for them. As with Bollinger Bands, the uses of STARC Bands vary from using an approach of the bands as a mean-reverting signal to some people viewing a breakout of the bands as an indication (as with moving average envelopes) that the price will continue strongly in that direction.

For the purposes of trading regime analysis, however, it is not the analysis of what the price subsequently does as it reaches the level of the bands that interests us. What the price does at these band levels is of little relevance to the trading regime analyst because we are trying to gauge just the expansion and contraction of the price action. We are interested in identifying the likely trading regime and so, like the Bollinger Bands, it is the width of the STARC Bands that will be of most interest to us.

Figure 15.5 shows the exchange rate of Japanese yen to US dollars in the daily time fractal and a STARC Band Width chart below using a 14-day ATR and a 10-day simple moving average for the STARC Bands. We can see that price range expansion and contraction does seem to be coincident with the expansion and contraction of the STARC Band Width. As with the Bollinger Band Width it is the direction and the extreme readings that interest us the most. When the band width is low it tends to lead to a subsequent period of trending, and when it is high it can lead to trend exhaustion and potentially sideways range trading. Although this chart shows a short-term time fractal the STARC logic works just as well on longer time fractals giving the analyst another tool to assess the long-term trading regime of markets.




Figure 15.5 USDJPY daily and STARC Bands

SUMMARY

In this chapter we have examined the role that discontinuous trend-following systems, such as the Turtle system, can play in determining whether a market is in a trending regime or a non-trending regime. We have seen that, as with all trend-following methods, there are drawbacks in the fact that it is non-predictive and lags the price. However, the occasions when there is persistence in the trading regime more than make up for the occasions when there is no persistence.

16

A Nod to the Quants

MARKOV CHAINS

Quants, eggheads, rocket scientists. Those who know them love them. These are the quantitative analysts, econometricians and traders who are now a common part of any trading or investment concern and tend to be the folk who sit in the corner of the room surrounded by whiteboards upon which are masses of algebraic equations, while having hugely animated discussions over conditional probability or space state! They are usually in the office before everyone else and leave late. They are driven people who are passionate about applied mathematics.

Applying quantitative techniques to financial markets has been a well-established trend for a number of decades now but the sophistication of the techniques has grown rapidly over recent years in conjunction with the growth of the hedge fund sector. Hedge funds and absolute return funds are “alpha seekers” many of whom search into the minutiae of financial market data in order to come up with an edge where they can extract value (money) from the markets. Lots of this analysis is done to look for arbitrage opportunities, where instruments like options or indices are mis-priced relative to the underlying assets or markets, but a great deal of quantitative analysis is still focused on anticipating the direction of markets; that is, the analysis done for trend-following models, range-trading models, neural network models and genetic algorithmic models (that “self learn” like some sort of Frankenstein monster!).

There are many quantitative teams that work in banks, hedge funds and investment managers around the world, but of the many I have come across only

a few stand out as being exceptional. I would rate the current (2008) team at The Royal Bank of Scotland headed by Pete Eggleston as one of those exceptional teams because they have managed to blend the required technical ability of physics, mathematics and general “rocket science” with an appreciation of the issues that investors and traders face in the markets in order to come up with straightforward and understandable quantitative solutions or alternatives. The RBS team have been working on the question of trading regime (in the sense of trending versus non-trending regime) for a number of years and have taken a quantitative approach, with the basis of that being an application of Markov chain analysis.

Now, this book is much more about the philosophy of trading regime as well as softer, more subjective forms of trading regime analysis, and so, while I could spend a long time going through the theory of Markov chains and how they are applicable to the financial markets, that is really beyond the scope of this work. Instead of blasting the reader with pages of calculus and equations, I aim to give a flavour of the thinking and analysis behind the method with the transcript of an interview with the RBS team as well as some very broad background in language that is hopefully understandable to other non-physicists like me. I would urge readers who are interested in this quantitative area of trading regime analysis to contact the RBS team for more information.

Named after the Russian mathematician Andrey Markov (1856–1922), a Markov chain is a so-called discrete time stochastic process that has the Markov property. The Markov property refers to the assumption that, given the present state (or regime for the purposes of our analysis), future states are independent of the past states. Put another way, the present state description fully captures all the information that can influence the future evolution of the process. Thus, given the present, the future is conditionally independent of the past, meaning that nothing that has happened in the past can influence or determine the outcome in the future. When we toss a coin, the chances of a head or a tail occurring is always 50% no matter whether we have just tossed 1,000 tails in succession. The idea is that the past outcomes do not influence the future outcomes.

Being *conditionally independent* means that the independent probability of events happening should only be influenced by another event happening. A classic example might be to think of two employees of a company, Robert and Bruce, who each live on opposite sides of a city. Robert gets the train to work while Bruce drives his car. Intuitively we would think that the probabilities of Robert being late for work and Bruce being late for work on any particular day are independent of each other. However, what if there was a train strike? All of a sudden the probability of Robert being late for work is affected (by the lack of trains) but the probability of Bruce being late for work is also affected (by the fact that the lack of trains leads to extra traffic on the road). The probability of Robert being late and Bruce being late are conditionally independent, given the probability of a train strike.

The theory of conditional independence and Markov chains can be found in many different real-world applications. When we hear the recorded telephone message “your call is important to us, please hold” the “queuing theory” behind this process is based on Markov chains. The importance or ranking of web pages on the internet as defined by search engines are driven by Markov chains, as are a lot of speech recognition models. In gambling, or games of chance, Markov chains can feature heavily and also where algorithms (or decision procedures) are used to create music.

Markov chains are all around us and are the bedrock upon which the RBS trading regime model is built. The following is a transcript of an interview conducted by the author (MG) with Pete Eggleston (PE) and Theodore Chen (TC) from the Royal Bank of Scotland about their model and its application. This was conducted by telephone in early 2008.

MG: Hi Pete and Theo, could you expand on the thinking over how the project came about, the genesis of the project for the Trading Regime model that you have designed, and any thoughts why Trading Regime, in your view, is important?

PE: We started the work back in 2002/2003, and the reason that we went down this route is that we were talking to a number of investors who were trading momentum models. Some of them had experienced some drawdowns and they were looking to try to improve their momentum models. So they were talking to a number of their banking counterparties on ideas for trying to improve these models, but our view was that you shouldn't necessarily be going to try to find what the next best momentum indicator is. Our view as a strategy was that momentum clearly does work and the way to try to improve it was to try to figure out when it works and when it doesn't seem to work. We were looking at market behaviour and really recognising the patterns where you clearly do see markets existing in different regimes. There was a very clear picture that at times there is a very strong uptrend, there are strong ranges, there are very strong downtrends as well as being periods where the markets are in high-volatility regimes and low-volatility regimes. But our key was that we wouldn't go and investigate or research, what the next best momentum indicator would be. Instead, we tried to start looking at ideas to try to identify when to trade momentum and when not to trade it, and that's what got us down the road of looking at regimes and really the Markov chain idea was one of the first we turned to. So it's not that we looked at lots of different areas here, it's just that we kicked off with it and it seemed to work when we started putting the thing together to try to calibrate it to try to identify these regimes. It seemed to look quite good. At the outset, one of the issues is that you have to decide how many regimes there are – so there's a certain issue of predetermining, if you like, how we think the market behaves. But allowing for that, basically we thought that within this [Markov] framework it looked quite encouraging. So that was really how it all kicked off. It was really just trying to improve the momentum trading model, although it mushroomed into many more applications.

MG: That's interesting in terms of the number of trading regimes, because in a lot of the previous work on regime, particularly in relation to Markowitz* optimisation there are multiple regimes. So you have approached it with a two-state environment where you wanted to look at whether there was a trending regime or a ranging regime?

PE: Well there are two components to what we have in this Markov process. One of those is based on looking at price and there we have three regimes – it's an uptrend, it's a range and it's a downtrend. And then, separate to that but all within the same framework, there's a two-stage framework that's looking at just the volatility, and we're looking there at high and low volatility. The reason we chose two states of volatility is that it had basically been supported anecdotally by our options traders who felt that vols existed in one of two states. We also did some work quite a long time ago, with King's College in London, looking at some particle filtering of tick data and that also seemed to empirically indicate that ninety-nine percent of the time currencies existed in two different volatility states. So that's why we chose the two-state vol model. The three states (of price) was really because we were trying to identify the trading signals of strong trends and so we were looking at, obviously, when to buy and when to sell. We did impose a ranging regime because we wanted to try to avoid those choppy ranging periods, so that was more of a pragmatic decision rather than based on any empirical evidence.

MG: Right, so essentially you have these three states of regime and two states of volatility, but how does the actual underlying details of the model work in that regard? How does the model come up with a decision that it is now probably a trending regime or is now probably a ranging regime?

PE: Basically what it's doing is using historical data, as obviously all these things do. So, taking a time series of historical data it is basically calibrating to a period of historical data to establish within these regimes what the means and standard deviations are in each of those regimes, and also what are called the transition probabilities; the probabilities that if you are in an uptrend today, what are the probabilities of being in the range tomorrow, or being in a downtrend, or still being in an uptrend? Now, what the model is doing is using this period of historical data to figure out these probabilities of moving from one regime to another. I don't know if Theo wants to add any more to that.

TC: Yes, we start off with a thousand days of historical data, which forms our sample set. We impose boundaries for each regime and take samples of means and standard deviations of each regime within these boundaries. We then

* Harry Markowitz, the economist who is best known for his work in studying the effects of asset risk, correlation and diversification on expected investment portfolio returns. His work is the basis of the Capital Asset Pricing Model (CAPM). A lot of work has been done in relation to Markowitz portfolios and "regime switching" but the "regimes" referred to on most of this work are mostly identified from a macro-economic, econometric and business cycle perspective rather than a simple volatility perspective.

construct probability distributions of each regime, based on these sampled means and standard deviations. From mapping our time series of historical data to these distributions, we obtain the time series of state probabilities. Obtaining the state probabilities is not enough; we then need to take samples of the states themselves based on these state probabilities. In other words, we randomly generate returns or prices and see whether it falls into an up, range or down regime, so we end up with a time series that looks like, e.g. (1 2 2 3 2 2 ...) of length one thousand days. This is our Markov chain and we can work out the number of times, for example, uptrend goes to range. This forms our transitional probability that then loops back to update the accuracy of our state probabilities. We repeat this whole procedure many times in a Monte Carlo simulation. Obtaining the means, standard deviations and transitional probabilities forms our calibration step. The state of the very latest data point will be considered as our current regime state. We can recalibrate every day as we add our last data point to our sample set.

MG: How did you decide on the length of the look back window?

PE: We've looked at lots of different things there, so when we first implemented the model we just had an expanding window and every new day we got a new set of data in and we'd recalibrate it on an expanding window. We have also looked at a moving window, and to be honest with you, for the majors [major currencies] it doesn't make a huge difference, the parameters are really quite stable. I think the key thing is that you need to make sure that the window is large enough in order for the thing to calibrate properly, so two years worth of data probably isn't enough, but four years would be. So whether you use a four-year expanding window or four-year rolling window, whatever you want to do is up to you. The testing we did on this last year doesn't seem that sensitive because the parameters are quite stable.

MG: So essentially the model has eventually one output saying that it's either going to be trending or ranging, is that right?

PE: The output is a matrix of probabilities, that basically gives you the probability today that EURUSD is either trending up, what's the probability of it ranging and what's the probability of it trending down? They obviously sum to one, and then separately you get the probability that EURUSD is in a low vol state today, or a high vol state today. They also separately sum to one.

MG: So how do you combine both of those outputs?

PE: The way that we would then construct a trading signal would be really very simplistically, and we haven't done any work to try to do this in a more systematic fashion because clearly there is a relationship, we think, between volatility and momentum. At the moment, the way that we have done this is that we say that if the probability of being in an uptrend is greater than ninety percent and the probability of being in a low volatility environment is greater than ninety percent, then we will initiate the trade. We would buy. And, basically, that buy signal would remain until either one of those probabilities fell below ninety

we're going to be looking at more of those futures markets, so I think that's very applicable. With regards to the options market, we have also published a paper on this towards the end of last year, using this information of these different regimes to try to trade volatility. And that looks very encouraging because it seems to make money in its own right after transaction costs, so that's great, but what's even more encouraging for me is the fact that when you look at the results of the vol strategy combined with the cash strategy there is a lot of very nice diversification. So our trend model had a pretty quiet 2006, for example, but if you look at trading vol in 2006 using this regime model, it did very well. Therefore, looking at it on a portfolio basis it starts to look very interesting as a total strategy. So we're looking to basically buy or sell vol depending on what sort of vol regime we feel we're in. And lastly, the work that we've just been doing is looking at, when this model says you're in a range, could we then do something along the lines of, say, double no touch barrier options or range accrual type products. I am just back testing that at the moment and that also looks very encouraging. We basically buy a double barrier or a double no touch, for example on knock out on, say plus or minus five percent around spot, and where this model says ninety percent or greater probability of being in a range. We close that out when that range probability decreases, ignoring the vol information at this point, and that looks pretty good as well.

MG: That sounds exciting. One last question perhaps is that, apart from the things you've just mentioned there, do you think there is any further future uses of trading regime models?

PE: Yes, I think trying to get a nice diversified portfolio strategy up and running is critical, because obviously the way we have this configured, there are long periods where the FX regime model doesn't trade. So, for example, so far in 2008 we haven't generated a trading signal and we want to obviously try to stay invested; therefore I think it's critical to try to have a variety of strategies available to you and this comes back to the concept of trying to figure out what works, what types of strategies work in what type of market conditions. So that's why we've been focusing on two different avenues – using all those probabilities to try to trade volatility, and also doing what Theo's been doing and extending into other asset types like interest rate futures. I think that is quite key. There are a couple of other areas, I think, that are potentially interesting and one came about because, at the moment we do all this in a univariate state; we basically look at EURUSD, and we look at whether EURUSD is going up, ranging, or down and whether it's in a high- or low-vol state. We started to do a little bit of work last year with Oxford University on what's called Coupled Markov models to try to look at the dependency of the different regimes; that is, if EURUSD is going up in an uptrend, typically does the USDJPY also behave in a similar way, or not. So whereas people will generally just use dependency measures like correlation, we were thinking about looking at the dependency between markets actually in regime or state space. So that's something we're

percent. If that were to happen the trade would be closed down. So basically, what we're trying to do is identify the strong probability of trends with the very strong probability that you're in a low volatility environment.

MG: Right, but the probability of the low volatility environment is separate from the calculation of the probability that, say, EURUSD might be going sideways tomorrow.

PE: Yes, it's totally separate.

MG: And the probability of the EURUSD either going sideways tomorrow, up or down is based on the Monte Carlo simulation over the last thousand days?

PE: Yes, that's correct, yes.

MG: OK thanks. So is this the volatility filter that you talked about in your paper? That is, the high-frequency volatility regime information as a filter?

PE: Yes, for that vol information, it's probably just worth noting that it seems to be critical in terms of trying to improve the performance of this model and last year [2007] was a great example of that. So what we typically find is that towards the end of a trend we see that intra-day volatility can often increase. A lot of people obviously use volatility filters but they often do it using either the daily close data or implied vol data and so on. That doesn't tend to pick up too much of an early warning as regards to the increase in intra-day volatility, and this vol filter we use is calculated on a twenty-four hour volatility number based on thirty minute sample data; so it's the high-frequency nature of that which I think is critical in getting us out of the trade prior to the trend turning and the draw downs kicking in. Last summer is probably one of the best examples of that. We were long of currencies like Kiwi and Aussie and this high vol signal kicked in on 25/26 July and closed us out of those trades just before the turmoil that then started through into August [both currencies fell well over 10% in only three weeks]. And interestingly, if you look at implied vol, although implied vol did move slightly on those days they didn't move enormously and filters probably wouldn't have picked them up. So implied only really moved significantly in mid-August and then it would have been way too late.

MG: That was good. So the volatility that you're using is realised volatility, yes?

PE: Yes, realised vol using thirty-minute data and that seemed to be critical.

MG: Well that's interesting. It definitely backs up what a lot of orthodox technical analysis would say about ranges getting higher towards the end of trends, up to trend exhaustion and volatility blows off a little bit, so that's certainly some overlap. Maybe just a quick thought on using the regime-switching model in terms of practicalities. Do you think that, not just the RBS regime switching model, but say the other regime type models or methods are best used on, for example, the cash markets or the options market?

PE: Yes, interestingly I think the directional signals are appropriate to anything reasonably liquid and we have applied it to lots of different markets so we have obviously applied it to spot FX, both in developed and emerging currencies. Theodore has been very busy recently in applying it to interest rate futures, and

quite interested in as well, and ultimately it's a more multivariate case rather than our current univariate case. So that's pretty complex and I think there's an awful lot to do there but I think that's potentially quite interesting. Connected to that is this whole idea of trying to connect the vol state and the changing state, and that's something that we're keen to work on as well. Also, what Theo's been looking at is actually modelling implied vol using the regime model.

MG: All right, so taking the process you described, it's not really a GARCH process, more a process of whether vol is going to be up, down or ranging the next period, is that correct?

PE: Yes. And how does that look Theo, does that calibrate?

TC: It does calibrate, yes. I haven't back tested it yet but obviously we need to adjust the threshold of when we trade and when we don't trade. Or sometimes we don't even need to use the three state, we might just go with the two state and see how it goes.

PE: Another point there is that, and this is something else that we have been working on, looking at correlation because we're also very interested in that. We know, for example, there are periods where correlations break down or when, for example, markets all start to move together and remain very correlated, and yet there are other times when there's a lack of correlation and so on. Again the question we have is, is there a defined sort of correlation regime in the market that we can try to identify as well? That's something we're keen to look at. Again, its early days and complex but I think that it could be quite an interesting way to try to model that sort of phenomenon as well.

MG: Well, best of luck with that. A lot of what you are describing, especially in relation to using different time frames for looking at the volatility change or trading regime, is similar to some of the more orthodox technical analysis techniques, so that is encouraging as there are obviously some crossover points here. Thank you so much for your time. It has been very informative.

Figure 16.1 shows the RBS Regime Switching Model in operation. The graph shows the chart of the euro to the US dollar (EURUSD) exchange rate using the left-hand scale while the right-hand scale refers to the series of solid and dashed grey lines ranging from 0 to 100%. What this is referring to is the probability of the market being in either an uptrending regime (the solid grey line) or a downtrending regime (the dashed grey line). When either line is at 100% the market should be experiencing the relevant trading regime. Conversely, in periods when neither line is showing a high probability of an up or down trending regime, then the conclusion is that the trading state of the market is a range-trading regime.

We can see from the chart that the model does a good job in identifying those periods when there is a strong, persistent up or down trend in the market and also, as in the middle of 2006, when there is a period of extreme range trading.

One of the most interesting aspects to this work is it backs up the case of a multi-time fractal approach adding value. The RBS model utilises high-frequency volatility




Figure 16.1 RBS Regime Switching Model Based on Markov Chains

data and allies that to a lower frequency model, and this is the same logic of analysing the trading regime in one time fractal but actually putting commitments into the market using a lower time fractal.

HURST EXPONENT

One of the most exciting developments in science over the last 50 years has been the emergence of Chaos Theory. This theory examines the observations that there is an underlying natural order to what would appear to be at first glance random data, and that this order is manifested in the emergence of fractal patterns. In this respect it is very similar to the Elliott Wave Principle and, in fact, there is some debate as to whether credit should be given to Elliott rather than the leading thinkers of modern Chaos Theory over the discovery of fractals in relation to the financial markets. One of the leading proponents of Chaos Theory is the celebrated scientist Benoit Mandelbrot, and the record shows that when he wrote an article for *Scientific American* in 1999 discussing Chaos Theory and time fractal analysis in relation to financial markets, he used (without reference) a very similar stylised chart to the one shown in Figure 8.1 in Chapter 8 that Elliott had first used in the 1930s. Could this be a coincidence? Perhaps, but notwithstanding the similarities between Elliott Wave and Chaos, the latter theory has provided some interesting observations that Elliott did not cover.

Chaos Theory has its modern origins in experiments conducted by meteorologist Edward Lorenz in 1960 when he was working on the problem of weather prediction,

the fact that it is persisting to be in a range of prices. Equally, what looks to the eye like a trend could, under Hurst, be deemed to be anti-persistent because the data is not showing much relationship.

Nevertheless, in reality H does a decent job in identifying where trends and range-trading regimes exist. We can see from the chart that when the EURUSD exchange rate was trading in its long-term symmetrical triangle range-trading regime in 2000, 2001 and into 2002, the Hurst Exponent was generally below 0.50. Also, when the market was experiencing trending regimes in 1999 into 2000 as well as 2002 into 2003, the Hurst Exponent was generally above 0.50.

This evidence then suggests that the Hurst Exponent could be a useful addition to our armoury of weapons to use when identifying the likely trading regime of the market.

SUMMARY

This discussion completes the section on existing trading regime analyses and techniques. In Part III we shall look at some further ideas for examining trading regime and how they might help the analysts or traders in their overall investment processes.

Part III

Further Ideas for Trading Regime Analysis

“Playing the game I have learned the meaning of humility. It has given me an understanding of the futility of human effort.”

Abba Eban (on golf)

but one of the backbones of the theory was developed by a man called Harold Edwin Hurst in the first half of the twentieth century. Hurst was a hydrologist who was working on the River Nile dam project. He was concerned with the problem of reservoir control and the question of what particular water discharge policy should be set so that the reservoir never overflowed or never emptied. The assumption Hurst started with was that rainfall follows a random walk but having tested his model by, among other things, dividing the range of the fluctuation of the reservoir level around its average by the standard deviation of the observations, he concluded that natural systems follow a biased random walk or a “trend with noise”. The Hurst exponent (H) that falls out of his model can be used to measure the strength of the trend and the level of noise, so it can be applied to any time series in order to examine the level of persistence or anti-persistence in the data.




Figure 16.2 EURUSD exchange rate 1998–2005. (Source: Bloomberg L.P. © 2008 Bloomberg L.P. All rights reserved. Reprinted with permission.)

Figure 16.2 shows the EURUSD exchange rate using the weekly time fractal from 1998 until 2005. The middle chart shows the calculated Hurst Exponent using a look back period of 25 weeks and when it is above 0.50 the data is said to be persistent, but when it is below 0.50 the data is said to be anti-persistent. Now, persistence or anti-persistence in this sense is not exactly the same as trending or range trading because, in the strict definition and calculation of H, a data series could exhibit persistence by

17

Implied Volatility Curves

In Chapter 3 we examined the role of volatility and discussed the merits of implied volatility (the volatility implied via options pricing) in the overall structure of trading regime analysis. We found that both historical volatility and implied volatility, although being good measures of the standard deviation of the rate of changes in prices, provided mixed performance when it came to using them for identifying trading regimes. However, volatility is not a one size fits all measure. We can calculate the historical volatility of the last N periods, where N is a variable number and this obviously affects the measures that we get. Implied volatility on the other hand looks forward not back, and when we want to know the level at which implied volatility is trading we need to know, firstly, the time period or “tenor” we are talking about. Implied volatility derives from option pricing and the option contract – an instrument that gives the buyer the right, but not the obligation, to buy or sell the underlying asset at some point in the future – will have to have a definitive expiry date. It is these expiry dates that become a major influence on the pricing of an option and thus, including over-the-counter (OTC) and exchange-traded options, we have implied volatility levels for one week, two weeks, one month, three months, six months, etc., and any “broken dates” within these tenors. The different levels of implied volatility for the different tenors make up an implied volatility curve not unlike a yield curve, and this enables analysis to be done on the curve too.

The expiry date of an option is vital to how much an option might be worth or what it is priced at. The logic of pricing is simple enough. If the option expires next week, then it will generally be worth less than the same option with an expiry date of three months away. This is the time value of an option, and the time value will be worth more the longer the option has to expiry because very simply it means that there is more time for something to happen that will affect the value of the underlying asset. Generally, therefore, we would expect the market to price the chances of something happening further out in time a little bit higher than something happening nearer to the present. So an option’s time value is dependent upon the length of time

remaining to exercise the option as well as the volatility of the underlying security's market price. The time value of an option decreases as its expiration date approaches and becomes worthless after that date. This phenomenon is known as time decay (or theta) and therefore anyone who holds options is holding wasting assets. For in-the-money options, time value can be calculated by subtracting the intrinsic value from the option price. Time value decreases as the option goes deeper into the money, whereas for out-of-the-money options, since there is zero intrinsic value, the option price will be all the time value. In general, time value increases as the uncertainty of the option's value at expiry increases, and this means that as the time to expiration grows the implied volatility of an option will, under normal market conditions, be slightly higher. In this way we can build, much like interest rates, an implied volatility curve.

Once we have our implied volatility curve we can use this to make inferences regarding how the market is thinking about risk. If the market is in a normal situation we can expect the implied volatility curve to be "normal" with an upward slope from the near dates to the far dates. In this situation the market is reflecting a psychology that everything is as it should be, with implied volatility for the near term being lower than implied volatility for the longer term. This is a normal situation because the more time there is for something to happen then the more time there is for uncertainty and, therefore, the higher the implied volatility should be. However, what if the implied volatility curve is showing a situation whereby the near-dated implied volatility levels are higher than the far-dated levels? This would be a situation where the implied volatility curve is said to be inverted, and the psychology of the market here is one in which there is a great deal of nervousness or fear that volatility of the underlying instrument is going to be greater in the near future than in the longer term. In this sense the market is not in a "normal" situation and the psychology of the market is in some sort of extreme state.

The same applies to the other extreme – that is, if the market is not just in a normal shape but that the near-dated implied volatility levels are *much* lower than the far-dated levels. In this sense the curve would be described as being extremely steep and the psychology that is being reflected in this pricing of implied volatility is one of extreme complacency and distinct lack of nervousness over the short term versus the long term.

Figure 17.1 shows three different positions of the implied volatility curve and the psychology that the pricing implies. As trading regime analysts it is this implied psychology that we are interested in because, if the options market is reflecting a certain psychology then that psychology will impact the pricing and trading of the underlying cash market, and we know that the trading regime of the cash market will be directly related to the current psychology of the market. The logic of the implied volatility curve psychology and its impact on the underlying cash market is straightforward. If market participants are extremely nervous we can assume that they will not want to put on big directional positions, one way or the other, and if that is the case, in all probability supply and demand in the market will generally become




Figure 17.1 Implied volatility/time graph

more evenly matched (in equilibrium) and a range-trading regime should probably be anticipated. Conversely, if market participants are extremely complacent we can assume that they will have been lightly committed to the market as prices have not really moved much and, if that is the case, in all probability supply and demand in the market will be in danger of moving from equilibrium to disequilibrium, and a trending regime should probably be anticipated. This is a logical way to think about how a market's overall psychology will be affecting behaviour and thus prices, but does it work in practice? By studying market psychology we should be able to get a good idea of when the market is extremely relaxed and when the market is extremely excited.

Zones of extreme relaxation tend to be followed by trends because the market participants are probably lightly committed to the market (do not have any large positions or bets) and so when the market starts to move they will tend to "wake up" to the market activity and get involved one way or another. Whether we like it or not this is how human beings behave. When the herd starts moving it is very difficult to be the lone wildebeest out there on the savannah, all alone! The natural animal urge within us is to follow the herd and by doing so we create trends in the financial and commodity markets. I can hear the fundamentalists among you protesting that you would not be influenced by the price action and would only decide to invest or put a position on in a market after you had well thought out, thorough and logical reasons to invest. Maybe that is the case for some people but they are generally the last people into the trend because only by then have the "reasons" to invest in the market come to light. I am not saying that everyone in the market is some sort of zombie-like trend follower, but what I am saying is that when prices start to move off on a trend it becomes increasingly difficult for market participants to fight that trend.

Zones of extreme excitement or nervousness on the other hand tend to be followed by periods of calm in the market. Excitement reaches a crescendo at the zenith of a trend, but excitement is a double-edged sword. For everyone who is long of a market

that is going up there will be an equal and opposite position (not numbers of people but position) that is short and vice versa for a market that is going down. Therefore, for a market that is going up in a strong trend there will be one crowd (the bulls) who are feeling increasing joy and euphoria whereas their counterparts (the bears) who are short of the market, either by being physically short or by not being invested, will be feeling increasing sadness and despair. At some point, these feelings will become so great that they will overwhelm each set of participants to the extent that the bulls cash in their chips and/or the bears throw in the towel and buy the market.

This is one reason for getting large volume spikes in the market at the exhaustion points of trends. These are the climactic moments of fear and greed when the pressure on those involved in the market either to close out their winning or losing trades is so overwhelming that all logic or reason goes out of the window. I like to call these moments POMBuCs which stands for Points of Maximum Buttock Clench. Graphic and crude it may be but I'm sure the feeling it describes will be familiar to most market participants. You know, that feeling you have when you realise you have just made a huge mistake or, alternatively, when you think you have just won the lottery? A real, stomach-turning moment.

In his book *The Day Trader*, Chicago Mercantile Exchange legend Lewis Borsellino tells a great story of one such moment. Borsellino was a futures trader in the S&P index pit and in 1987 was on holiday in Europe at the time of the stock market crash. As soon as he heard what the markets were doing on the Monday of the crash he flew back to Chicago, arriving at the "Merc" around the middle of the trading session on what became known as Black Tuesday. He didn't waste any time in going to work in the pit and made \$500,000 that day and the next as the market bounced a little. On Thursday morning, he tells of how he felt something in the air before the opening of the S&P. The order clerks were jumpy and he surmised that there must be heavy sell orders that big institutional customers were wanting filled at the opening. Just before the opening bell traders in the open outcry pit shout out bids or offers to gauge where the market is going to open relative to the previous close. Borsellino locked horns with a broker from an institution who was offering to sell "1,000 (points) lower". Borsellino offered "2,000 lower" and this went back and forth with Borsellino lowering the opening indication every time. When the opening bell rang the S&P future had opened 5,600 points lower and Borsellino (who had played the main part in offering the market down pre-open) then bought 150 futures contracts. Less than a minute later, he sold those contracts at 2,000 points higher than his purchase price making \$1.3 million for his own account. He was 30 years old and had made \$1.3 million in less than a minute. He writes, "I handed my last trading card to the clerk, walked out of the pit, went into the bathroom, and threw up . . . I tasted the bile in the back of my throat and knew what my gut was telling me – it could have just as easily gone the other way . . . If the market had turned against me, I would have been wiped out."

So call it puke point, buttock clench or whatever you like. The point is that market participants feel these emotional extremes and it is these emotions that dictate our

actions in the market. In fact, the people doing the most puking that day were the people who *sold* at the opening. They were selling *after* a big move down in the market and this point of maximum panic is what Borsellino identified and took advantage of by going the other way.

These climaxes of emotion are just that. Most market participants are exhausted after all the mental and, in the case of open outcry, physical stress involved in seeing their positions either climbing in value to the point at which they are seeing dollar signs in their eyes like cartoon characters, or losing value to the extent that the only picture they have in their heads is of them and their family in a gutter with a sign saying “brother can you spare a dime”. After these moments it is hardly surprising that the markets will tend to enter into quieter, calmer periods.

Figure 17.2 shows the line chart of the EURUSD exchange rate from 1999 using the weekly time fractal. This is the left-hand scale. Overlaid on the chart and referencing the right-hand scale is the difference between the three-month and one-week implied volatility rates for the EURUSD exchange rate and, because this analysis is measuring the state of nervousness in the market, I have called it a Nerve Index. When the line is below zero the one-week rate is below the three-month rate and the volatility curve is normalised, whereas when the line is above zero the one-week rate is above the three-month rate and the curve is inverted. Therefore, when the line is below zero we can conclude that market participants psychology is relaxed, but when the line is above zero we can conclude that market players are feeling a bit nervous. What the index is capturing is the market pricing of very short-term psychology




Figure 17.2 EURUSD Nerve Index weekly

relative to slightly longer-term psychology. When the volatility curve is inverted the market is demonstrating a fear of the market moving a lot over the next week relative to what is expected on a three-month basis. In this way it captures a dynamic of relative nervousness and, following the logic of trading regime analysis, what we would expect to see is that, at the extremities, there could be turning points in trading regime.

The chart shows that when the index displays extreme readings it can, in fact, coincide with a regime switch. For instance, the EURUSD had been in a long, persistent downtrend since the euro's launch in 1999 and we can see in 2000 the implied volatility index spiked up on a couple of occasions, indicating extreme nervousness in the market. On each occasion it precipitated at least a short-term pause in the trend but the final spike up in September 2000 came almost right at the bottom of the long trend (when emotions were particularly extreme that the EUR was in freefall and nothing could stop its fall!) and this emotional extreme ushered in a regime switch and a long period of range trading.

This period of consolidation turned out to be a symmetrical triangle and we can see that during that time the Nerve Index was generally below the zero line, indicating that the market psychology was relatively relaxed. This is as we would expect during a period of consolidation because the pricing of implied volatility will generally be coincident with the underlying price action. As the market goes sideways, option market makers factor in that this sideways trend will continue into the future, and so implied volatility comes down. In this case, the short-term volatility comes down relative to the longer-term volatility and the volatility curve becomes normalised. Towards the end of the symmetrical triangle, in November 2001, the Nerve Index reading spiked lower to below 3, meaning that the one-week implied volatility was three "vol" below 3-month, the implied volatility.

Clearly the market was pricing that there was nothing to worry about short-term volatility relative to longer-term volatility, but the extent of the difference between the two volatility levels was a signal to the trading regime analyst that there was an *extreme* level of complacency that the market was going to continue in this range bound trading regime. These are the times when the trading regime can switch and only a few weeks later the market broke out of its consolidation into a strong trend. As this trend developed we can see that the implied volatility curve was becoming more and more inverted as the market became more fearful of short-term volatility relative to longer-term volatility.

At the end of June 2002 the EURUSD had moved up close to the parity level and this coincided with an aggressive move up in the Nerve Index, indicating an extreme move in the level of nervousness in the market. Again, we can observe that this extreme moment came close to the time that the trend paused and the market went in to a period of range trading until November of that year.

The EURUSD continued on its uptrend with another trend exhaustion climax in June 2003 but, in December 2003, something odd happened. We can see that the price of the market was trending higher up past 1.20 and the Nerve Index was relatively

steady until the end of December when it suddenly spiked lower, only to spike higher immediately afterwards. If the logic was being followed, the analyst will be asking why the volatility curve spiked lower when, with the trend ramping up quite hard, it should have been expected to be rising above zero. Unfortunately this highlights one of the problems with using options volatility data in an analytical setting in that, on very rare occasions, the market can, due to poor liquidity, market positioning or a combination of both, act in a manner that is clearly at odds with logic. The month of December can be a time when liquidity in the markets is very poor and, at these times, the analyst must be aware that strange things can happen to prices, especially in options markets. However, although the index spiked lower, it did so only very briefly and within a couple of weeks had risen back up to the levels we would expect it to show at the climax of a trend.

As usual, the question of what constitutes an extreme level of this volatility curve measure is the one that is most difficult to answer. We can look at the historical highs and lows that the index has achieved and say that is the level of extremity. However, the market will not always go to these extreme levels when signalling a potential climactic moment in either nervousness or complacency. We can use normalising techniques to contain the upper and lower limits, but the issue does not go away. The answer comes down once again to using an analysis like this in conjunction with other analyses so that a case can be built for a regime switch by gathering many pieces of evidence.

SUMMARY

In this chapter we have taken a brief look at how the changing shape of the implied volatility curve of a market can be used to form an analysis of the potential trading regime of the underlying market. How the options market prices implied volatility across different elements of time tells us a lot about the underlying psychology of the market and whether it is skewed towards nervousness or relaxation. In the next chapter we will dig deeper into how the market prices implied volatility at one period of time and how that can help us with trading regime analysis.

18

The Volatility Smile

“Are you looking at me?”

Travis Bickle (Taxi Driver)

The options market is an extremely useful market through which one can analyse market sentiment. This is because the manner in which options are priced reflects the attitude that traders and investors have towards certain types of risk in the market. It is the actual pricing of risk that reflects the *real* sentiment of traders much better than a survey of traders' attitudes, because the pricing of risk contains real, live positions and potential positions whereas a survey of attitudes can reflect the fact that sometimes people say one thing and do another. Surveys can also include an element of hope and bias, whereas the price is unbiased.

One of the most important factors in pricing options is the level of volatility implied by the price of the option. Implied volatility reflects the market view of how volatile the market might be in the future, and it has a marked effect on option prices. If implied volatility is lower than normal, the options price will be lower than normal; and if implied volatility is higher than normal, the options price will be higher than normal, all else being equal. Option contracts contain strike prices that vary up and down from where the market is currently trading, and the further away the strike price is from the current underlying market rate the more the delta of the option (the probability of the option expiring being worth something) tends towards 0 or 1. That is, if an at-the-money option has a delta of 0.5 then an out-of-the-money option is going to have a delta below 0.5, with an in-the-money option having a delta above 0.5. The more the option is out of the money the more the delta tends towards 0 and the more in the money the option becomes, then the more the delta tends towards 1. At any one point of time there will therefore be various strike prices that could be dealt on and each one will have a slightly different delta depending on where it is in relation to the current underlying market. If we think of the range of delta possibilities as a strip, then it could look something like

Figure 18.1, which shows the range of option deltas from out the money to in the money. The top line shows the range from 0 to 1 reflecting the delta of the option depending on whether it is either a call option or a put option. The bottom row in italics reflects the fact that the delta on one side will be equal and opposite on the other. That is, a 25-delta EUR call option will be a 75-delta EUR put option, and vice versa.




Figure 18.1 Option delta range

PUT/CALL PARITY AND RISK REVERSALS

The definition of put/call parity is essentially the relationship between the price of a put option and a call option on a stock according to a standard model. Specifically, in layman’s terms a call option (to buy a stock) must be priced equally to a put option (to sell the stock) if the delta (the probability of the option expiring in the money) is the same for both. An intuitive way to think of put/call parity is that, given the fact that a long physical stock position could be synthetically created by buying a call and selling a put, if there was any mis-pricing between the options, then an arbitrage opportunity would exist for traders to exploit.

One option strategy that falls out of the theory of put/call parity is that of the “risk reversal”. The risk reversal option strategy involves the purchase of a call and the sale of a put with the same delta, and it is generally used as a cheaper hedging strategy compared to a plain vanilla call or put. As an example, let us suppose that an investor who owns IBM shares is concerned that they will fall in price. He could buy an out-of-the-money plain vanilla put option to hedge against this risk and this would cost him perhaps a 1% premium. However, rather than pay this premium he could sell an equivalent delta in-the-money call option for 1% (it is the same price due to put/call parity) and effectively finance his hedge for zero cost. So if IBM shares fall sharply, he has hedged his exposure and has done so at zero cost. Well, he has done so at zero cost on a nominal basis. The hidden cost would be that if IBM shares rallied hard above the sold call option strike, he is effectively locked into selling IBM there rather than the higher price that could be achieved in the cash market.

So the risk reversal option strategy is one that takes advantage of put/call parity to finance an option hedge at zero cost. However, in practice, it is not always zero cost and this is because in the real world of options trading the put/call parity theory has a little poetic licence.

SMILE PLEASE

Put/call parity sounds fine in practice. In reality it is somewhat different because the Black–Scholes option-pricing model has one fundamental flaw as it assumes a constant level of volatility across the delta spectrum. In practice, the market makers who have to make prices on options will price volatility a little bit differently according to the delta. Remember, the delta of an option can be thought of simply as the probability of the option ending up in the money, and so each delta level will mean something different to the person making the price (or deciding the level of volatility on which to price the option). If our market maker is pricing a 50-delta option (50% probability of ending in the money) he might price it at, for example, 10% volatility. However, if he was asked to price a 10-delta option his thought process might change to reflect the fact that, although there is a lower chance of the option ending in the money, if the market did move that far then the rate of change of that delta (the gamma) could be very fast indeed. In other words, although the 10-delta option has less chance of ending in the money, if it did end up in the money it has more perceived risk attached to it from the price maker's point of view. Our market maker might then price the 10-delta option with a volatility level of something like 10.3% to reflect this higher risk. That is, the 10-delta option is more expensive in volatility terms because it has a higher perceived risk even if it has less chance of ending in the money.

This perception of risk from the market may result in an at-the-money option that has a different level of implied volatility to an in-the-money option and an out-of-the-money option. The difference in the level of implied volatility can also be affected by the market's fear, for instance, that the market has more chance of falling than rising. In this case the puts will have a slightly higher level of implied volatility than the calls to reflect this heightened fear of the downside risk. Both of these attitudes towards risk by market makers create what are known as option skews, where puts and calls with the same delta may well be priced differently and it is this difference in implied volatility levels that create the volatility smile.

Figure 18.2 shows a stylised volatility smile in a certain tenor, for example one month. In this example the market in question is showing a very balanced position whereby the at-the-money options (those with a delta of 0.5) have an implied volatility of 8% and the out-of-the-money options (those with a delta of 0.1) have an implied volatility of 9%. From this data the analyst can infer that the market's pricing of the options is showing that there is a great deal of fear for both the upside and the downside. My nickname for this situation is Bonnie Langford, named after the British actress and entertainer who, perhaps unfairly, is renowned for having a big, perfect, permanent, toothy smile.

However, this balanced pricing of a Bonnie Langford is actually a very unusual situation. It is much more normal for the volatility smiles to be grumpier and/or skewed to one side or the other. This skewness reflects the fact that it is much more normal for the market to fear one side more than the other, especially when a big move or trend is happening.




Figure 18.2 The volatility smile

Figure 18.3 shows a situation where there is a skew in the options volatility strip (or surface as it is sometimes known) if there are multiple tenors to consider. The 50-delta at the money options are still being priced with 8% implied volatility, but the 10-delta calls are now priced at 8.5% whereas the 10-delta puts are priced at 9.5% implied volatility. This information is incredibly valuable for the analyst because it tells him where the options market fears the market move. That is, which direction of potential market move does the option market fear the most, or is most nervous about? In this case it is the downside because the put options are priced at a higher volatility than the equivalent delta call options and are therefore more expensive.




Figure 18.3 The volatility smile with puts bid

They are more expensive for a reason. If we think about it from a buyer's point of view, who is buying the option for insurance or hedging purposes, an option buyer will buy a call option if he wants to *protect* against an upside move in the underlying market and he will buy a put option to *protect* against a downside move in the underlying market. Essentially, in this example the market is telling us that downside protection is more expensive than upside protection because the market thinks there is more chance of a downside move than an upside move in the underlying market. This difference in price between what, on the face of it, are equivalent chances is something similar to a motor insurance company charging more for a policy to an 18-year-old male driver than to an 18-year-old female driver. Yes, they are both 18-year-old drivers and therefore inexperienced, but the male will have statistically more chance of crashing the car than the female due to the existence of quite large amounts of testosterone in his body that makes an 18-year-old male driver more dangerous than an 18-year-old female driver.

I call the volatility smile in Figure 18.3 the Mask in honour Jim Carrey's graphically enhanced exaggerated lopsided smile in the film *The Mask*. If it was skewed the other way it would show that calls were priced more expensively than puts and that the market was more fearful of an upside move than a downside move. We also have the De Niro (a horizontal line where all deltas have the same volatility in reference to the poker-faced actor Robert De Niro), a Dirty Harry (a horizontal line with a little blip up on the extreme 10 deltas in reference to Clint Eastwood's portrayal of the slightly deranged cop Harry Callahan) and a Gazza (the inverse of the Bonnie Langford in reference to the soccer player Paul Gascoigne who famously burst into tears during a World Cup game). The Gazza is extremely rare, the De Niro quite rare and the Dirty Harry appears more often. The normal situation, however, is a Mask type smile where there is some skewness within a distinct smile effect.

The serious point of all this flippancy is that by analysing the smile of the options market we can get a very good idea as to sentiment in the underlying market itself, and if we can do that, then maybe we can get a better idea of whether the psychology of the market lends itself more to a trending trading regime or to a range-trading regime. If the options smile is relatively flat (like a Dirty Harry or a De Niro) then the market is pricing the volatility in the wings (the outside deltas) at similar levels to the 50-delta at-the-money options. This pricing structure is telling the analyst that there is a lot of complacency in the market – that the market is just as nervous of a big move in either direction as it is about the market staying around the same levels. The market is pricing itself not to fear a trend in either direction and, therefore, to expect a trading range. Conversely, if the options smile is at a pronounced skew (the Mask) then the market is pricing the volatility in one direction as opposed to another. That is, the market is pricing one directional risk higher than the other which allows the analyst to conclude that the market fears a directional move or a trend in one particular direction, and the market is pricing itself to fear a trending regime. It is this logic of the psychology behind the pricing of options that could allow us to

construct an index of the volatility smile that should help us with another piece of the trading regime analysis jigsaw.

The Volatility Smile Index (VSI) is an indicator of the extent to which the options implied volatility strip is smiling at any one particular tenor. It seeks to neutralise the skewness in the smile and instead seeks to calculate just the extent at which the low delta options are being priced away from the at-the-money options. In this regard it eliminates any directional analysis from the calculation and is only concerned with whether the market fears a trending move in one direction or the other.

$$\text{Volatility Smile Index} = (10\text{-}\Delta \text{ call vol} + 10\text{-}\Delta \text{ put vol}) - (2 * \text{ATM vol})$$

The Volatility Smile Index adds together the 10-delta volatilities but you can construct any similar index using other deltas; the main point is that the deltas have to be the same. The philosophy behind the Volatility Smile Index is that if the market does not fear a big move in one direction or the other, then the volatility in the wings will be priced near the centre; but if a big move is feared, then the volatility in the wings will be far away from the centre. The important part to note is that one wing may well be skewed away from the other wing but if we only analysed *that skew* (between the wings), then it doesn't tell us how the wings are themselves priced to the centre.

For instance, let's consider a situation whereby the 10-delta puts are priced at 8%, the at-the-moneys are priced at 7.5% and the 10-delta calls are priced at 7.6%. There is a skew between the 10-delta options of 0.4% – the 10-delta puts are priced 0.4% higher than the 10-delta calls and therefore the market could be said to fearing a downside move in the market. Now consider the same market where the 10-delta puts are still priced at 8%, the 10-delta calls are still priced at 7.6% but now the at-the-moneys are priced at 7%. The skew between the 10-delta calls and puts has not changed and so the market still fears a downside move relative to an upside move to the same extent. However, the shape of the smile has changed dramatically; whereas before the central at-the-money options were priced at only 0.5% and 0.1% away from the wings (10 deltas), now the central at-the-moneys are priced at 1% and 0.6% away from the wings. In other words, whereas before the market was relatively relaxed about the extent of the *absolute* movement away from current rates the market is now much more fearful of a big move away from current rates. The market still fears a downside move relative to an upside move because the puts are still priced at the same relatively higher amount than the calls, but the market is now pricing a greater risk of a bigger move in either direction away from current rates.

It is therefore the information of how the market is pricing the movement away from current market rates (the at-the-moneys) that is important to the trading regime analyst rather than how the market is pricing the relative skewness between the lower delta options. As we shall see below, the skew analysis is still valuable and can be used for identifying extreme sentiment readings, and therefore potential trend exhaustion zones, but the logic of the volatility smile dictates that it is the smile

itself rather than the skew that should be of greater interest for an analysis of trading regime. But how does it work in practice?

Figure 18.4 shows the EURUSD exchange rate using the weekly time fractal and a Volatility Smile Index using the 10-delta wings with the index having been smoothed by using a five-period moving average. When the index is going down it indicates that the volatility smile is becoming less smiley and more “serious” as the aggregate 10-delta volatilities become closer to the adjusted at-the-money volatility, whereas when the index is going up the volatility structure is becoming more smiley. Using the philosophy described above we would want to see that a falling Volatility Smile Index coincides with a range-trading regime and a rising Volatility Smile Index coincides with a trending regime in the underlying market.




Figure 18.4 EURUSD and Volatility Smile Index

However, the evidence in this example does not show that a strong relationship exists to back up this philosophy. Yes, the index reached its lowest point and stayed there for a while as the market traded in a distinct range during the latter half of 2006, and the subsequent regime was a strong trend from 1.25 to just below 1.35. The index subsequently very generally ebbed and flowed with the market's regime cycles between ranging and trending, but the evidence is patchy at best. There is more research to be done here but, based on this evidence, the use of a Volatility Smile Index type of analysis will certainly enhance our overall trading regime analysis, but it should not be used as a stand-alone analysis in its own right.

rate. This makes intuitive sense because, as the underlying moves up, there would probably be increased demand for calls relative to puts and the option market makers would be hedging themselves by marking up the calls as well. It is this relationship that has led risk reversals to be used as an indicator of overbought or oversold levels (or trend exhaustion) and we can see from the chart that this philosophy can work at certain times. When EURGBP was in a range-trading regime the risk reversals would tend to rise or fall coincidentally with the underlying exchange rate in an oscillating fashion. However, once EURGBP started to move higher in a trending regime the risk reversals continued higher, moving well above the previous peaks. The old extremes of EURGBP calls being bid relative to puts were exceeded by a great amount, and so the value of the risk reversal as an overbought or oversold indicator suffers, like all oscillators do, when a persistent trend develops. Therefore, what does the trading regime analyst do with the risk reversal?

Perhaps a more subtle way to think of the risk reversal is like another oscillator we have looked at, the rate of change. When we use the rate of change as a trading regime indicator we essentially look for periods where the rate of change is slow (close to zero), indicating a range-trading regime or fast (an amount away from zero) indicating a trending regime. The risk reversal can be used in the same way but with a twist.

When the risk reversal is trading close to zero then the market pricing is reflecting a psychology whereby the risks are perceived to be balanced between an upside move or a downside move. In other words, the market is quite complacent and relaxed about a trending move occurring either way, and given this psychology the analyst would expect that the underlying market price would be somewhat range bound. On the other hand, when the risk reversal starts to move away from zero, the psychology of the market is becoming more nervous of a trend developing; but this is happening coincidentally with the move taking place. Therefore, if the risk reversal is trading close to zero we can say that a range-trading regime is probably the one to be favoured, but if the risk reversal starts to move away from zero we can say that a trending regime should probably be favoured. It is certainly not perfect by any means as we can see from Figure 18.5, but, as a general philosophy that the trading regime analyst can overlay on their other analysis, it makes intuitive sense.

The twist in the risk reversal story is that sometimes they can give the analyst a clue about the direction of a forthcoming trend. We have already seen that in the cash market it is possible for demand or supply dynamics to take place during a consolidation phase and this can give hints on the ultimate resolution of the consolidation such as in ascending or descending triangles, or by the volume analysis of a head and shoulders pattern where volume goes up on the down moves and down on the up moves. Similar forward-looking clues can be found in the risk reversals for the analyst who is prepared to look. Sometimes it is possible to observe changes in the risk reversal pricing that might suggest a building up of demand or supply dynamics during a quiet phase in the market that could give a hint as to the possible direction of the breakout into a trend. Looking at the EURGBP example in Figure 18.5, we can see

SKEWS

If a Volatility Smile Index cannot always help us with a reading on the potential trading regime, what about an analysis of the actual level of the risk reversal itself? The level of the risk reversal is different from the Volatility Smile Index because the index is measuring the “smiley” of the smile whereas the risk reversal is graphically displaying the skew. That is, if the volatility smile was a Bonnie Langford showing a symmetrical smile, then the index would show a high reading whereas, due to the fact that the low delta calls and puts were, although higher than the at-the-money options, both priced at the same volatility, the risk reversal would show no skew favouring calls or puts.

Figure 18.5 shows the EUR exchange rate with the GBP using the weekly time fractal and the risk reversal level of the one-month, 25-delta options.




Figure 18.5 EURGBP weekly and 25-delta risk reversal

The risk reversal level on the right-hand scale refers to the difference between the volatility levels for 25-delta calls and puts. Above the zero line indicates that EURGBP calls are being priced using a higher volatility than the puts, and vice versa below the line. Above zero, the market makers are therefore pricing options to reflect demand for calls relative to puts and/or they are nervous of an upside move in the market relative to a downside move. We can see from the chart that there tends to be a decent relationship between the risk reversal level and the underlying spot exchange

that from 2003 until early 2006 the risk reversal and the exchange rate were very broadly acting in tandem. The rate went up and calls became bid over puts and vice versa at the same time. However, in 2006 a subtle shift began to emerge whereby the underlying exchange rate was still generally range trading (and in fact falling ever so slightly) but, that during this phase, EURGBP call options were being more persistently priced higher than the corresponding 25-delta put options. That is, the risk reversal was more persistently above the zero line than before and, given the fact that the underlying exchange rate was still in a general range-trading regime, this observation could well have provided the analyst with a major hint that underlying demand for EURGBP was taking place, and that an underlying bullish sentiment was developing.

Eventually, as we can see, the market broke higher into a very strong and persistent uptrending regime, making the observation of this prior shift in risk reversal pricing quite timely. So examining the option skew data and how that is shifting in conjunction with the price movement can sometimes give a good lead to the analyst on underlying supply and demand in the market.

SUMMARY

In this section we have examined the volatility smile in the pricing of options and how that can help in an analysis of trading regime. Just as in the pricing of implied volatility at different tenors, the pricing of implied volatility at different option delta levels tells us a lot about what the market thinks in general about the outlook for the underlying price. Whilst it is not obvious that it can be used in any systematic manner, the shape of the volatility smile can give the analyst a good feel of the underlying sentiment in the market.

GAMMA AND RANGE TRADING

“Kurt, where is our jamma [sic]?”

The physical manifestations that a range-trading regime psychology has on the market is that some market players will try to play the range by buying if the market goes down and selling if the market goes up. In some respects this can add to the ability of the market to stay in its ranges. There is one sector of the market, however, where the behaviour is always to sell up moves in price and buy down moves in price, and sometimes if they dominate the market at certain times this behaviour can be a big factor in creating the range in the first place, or at least adding to it.

I am referring to those traders whose main activity is to trade in a “delta neutral” style and “scalp the gamma”. No, this is not referring to some sort of ancient native ritual but to traders and market makers who try to make a profit out of options

trading or combining options with cash positions. Financial markets are jammed full of jargon, which, to the layperson, seems to be baffling and esoteric and the options market in particular adds another level of jargon to this.

An option is a contract that is entered into between two parties that gives one party the right, but not the obligation, to buy or sell an underlying asset or market at an agreed price known as the strike price. The difference between the strike price and where the underlying market is trading is used to calculate the probability of whether the option contract will end up “in the money” or, put another way, whether it will be worth something when it expires. This probability of whether the option will expire with some worth is known as the delta of the option and the delta will move around during the option’s life depending on where the underlying market trades. The speed of the change in the delta will be affected by the speed of the change in the underlying market among other factors and so this, too, will be variable. This rate of change of the delta is known as the gamma of the option.

The delta therefore measures the probability of the option expiring with some worth and the gamma measures how fast the delta is likely to move. The gamma therefore becomes vitally important to people like market makers who want to remain as neutral as possible to underlying moves in the market. Remember a market maker’s job is, as the name suggests, to make a market, a price (a bid price and an offer price) to someone who wants to either buy or sell the market. They generally do not want to have a “position” in the market where a substantial move one way or the other will affect their “book” either negatively or positively. They leave that job to the proprietary traders who take directional risk with the bank or institutions capital.

As well as market makers, however, there are a host of other strategies that investment managers and, indeed, proprietary traders, use that results in the exposure to the market move (or put another way, to the move in the delta) being neutral. The term “delta neutral” means that the person who is holding the position wants it to be neutral to any movement in the market. At a first reading of that statement the reader might think it is an odd thing to do. Is not the whole point of investing or trading in the markets to take a position in the market? Well yes it is, but we must remember that the market can go essentially five ways from any point in time (up a lot, up a little, unchanged, down a little or down a lot). There are many strategies to take advantage of sideways moves in markets and being delta neutral is one method involved in these strategies. We must also remember that someone could have a directional position in the market but, for whatever reason, wants to hedge that position. He might be more uncertain as to the direction but doesn’t want to close the underlying position. Market makers will accumulate long and short positions all the time in their dealings but because of the nature of their business where they want to make small fractions of profit on a frequent basis (or indeed due to the mandate they are given by the bank or institution) they will hedge those long and short positions. They hedge these long or short cash positions with derivative positions such as options, and when they do that they will hedge with enough option contracts to be delta neutral on the overall position.

For example, let's say that a foreign exchange market maker has just made a EURUSD price to a fund manager and the fund manager has sold €20 million. This means that our market maker has just bought €20 million and sold the USD equivalent. He has accumulated a long position in EURUSD of €20 million. Now let's say that he wants to hedge this position and, furthermore, that he wants to hedge this position using the options market. There are various ways this could be done but let us assume for simplicity that he opts for the most straightforward method and buys a EUR put option. He holds a long EUR position in the cash market (he is physically long EUR) and so to hedge that he buys a EUR put option which gives him the right but not the obligation to sell EUR at a predetermined strike price. Note that buying the EUR put option will mean that the flip side to that on a foreign exchange deal means that he will have bought a USD call option. If the underlying exchange rate were EURJPY instead of EURUSD, then he would have bought a EUR put (JPY call) option.

So our market maker has decided to buy the EUR put (USD call) option and decides that he wants to buy an at-the-money option. This means that he will buy an option with a similar strike price to where the market is trading at the moment. If the EURUSD exchange rate is trading with a spot rate of 1.5617 and a one-month forward rate of 1.5591, then if the market maker wants to buy an option with a one-month expiry, the at-the-money strike price will be equivalent to the one-month forward rate of 1.5591. An out-of-the-money EUR put would be an option with a strike price below 1.5591 and an in-the-money EUR put option would be an option with a strike price above 1.5591. Due to the fact that an at-the-money option is, as the name suggests, struck with a strike price at the current market rate and, if we assume that the market has the same chance of going up as going down from any point in time, then what we are buying is an option that has a 50% chance of ending up in the money at the expiry date. That is, the option will have a delta (a probability of ending up being worth something) of 50. If the market maker had bought an out-of-the-money option then, as the name suggests, because it is out of the money, it will have a lower chance of expiring worth something and will therefore have a delta lower than 50. The opposite is true for in-the-money options where the delta (the probability of expiring worth something) will be higher than 50.

A cash position has a delta of 100 because it has a 100% chance of ending up either in or out of the money. The option that has been bought in this example has a delta of 50, meaning that it has a 50% chance of expiring in or out of the money. This means that the equivalent cash market position of the option will be half of the nominal option amount. If the option contract was struck in a nominal amount of €1 million, then, because it has only a 50% chance of expiring in the money, the equivalent cash position is 50% of €1m, which is €500,000. Therefore, in order for our market maker to hedge the long euro cash position and make the whole position delta neutral, he has to buy an at-the-money option in double the nominal amount of the long cash position. That is, an option with a nominal value of €40m will have to be bought to delta hedge a long cash position of €20m.

The market maker puts on this delta hedge but it is only a perfect delta hedge for the moment it is put on. Once the market starts moving (combined with the relentless march of time), the value of the option will change relative to the value of the long cash position, and it is this *change* in the options value relative to the long cash position that creates opportunities for the market maker to “scalp the gamma” of the position for profit.

By being long the option the market maker is considered to be long gamma. What this means is that he has a positive relationship with the rate of change of the delta of the option. When the market moves the delta of the option will move and this will, on a delta neutral basis, benefit the holder of the option. If he was short the option he would be short gamma and, when the market moved, he would, on a delta neutral basis, have a negative relationship with the rate of change of the delta.

Perhaps the easiest way to view this gamma effect is to think of a situation where you are either long an option or short an option with nothing else to worry about. So here we have the EURUSD exchange rate that is the underlying market and we can either buy an at-the-money one-month EUR put option or sell an at-the-money one-month EUR put option. Buying the option costs us money (premium) but that cost comes with protection as to our known potential losses. Selling the option will earn us money (premium) but earning that premium comes with some potential disastrous outcomes. Our market maker has then bought the put option and because this option is at-the-money it has a delta of 50 at the start of its life with a gamma of 14%. The strike price is 1.5591 and the current spot rate is 1.5617. The €40m nominal option hedge cost €444,092 in premium.

Now let's assume it is two weeks later and the spot market has moved down to 1.5200. The option now has a delta of 89 and is worth €1,108,079, so the trader has a profit of €663,987 on the option side of the equation. His loss on the euro cash side is €534,034 ($(€20m - €20m * (1.52/1.5617))$). The profit on the options side is assuming unchanged volatility and, in reality, it would have gone up, making the profit there even greater. Nonetheless, the trader now has a net profit on his hedge of €129,953. This has occurred because, as the market has fallen, the delta on the options has risen (at a gamma rate of 14%) whereas the delta on the cash side has remained fixed, and the option also has some time value left in it. The trader's position is no longer delta neutral and, therefore, to return to a delta neutral position he has to rebalance. He can do this by selling some of the put option he owns in order to realise a profit and get the position back to delta neutral. By doing this he is said to have “scalped the gamma” and this type of activity can, when it becomes a major factor in the markets at certain points of time, enhance a market's range-trading qualities. That is, if the markets were quiet and the option market makers were long gamma in general, then they will be one section of the market that will sell when the market moves up and buy when the market moves down, thus enhancing the ability of the market to stay in a range. It therefore makes sense to try to work out how the options market makers in a market are positioned and whether this sort of gamma-scalping activity will be prevalent or not.

SUMMARY

This section has discussed the impact that certain types of option-related trading can have on the trading regime of a market. The trading regime analyst should be cognizant of the impact that these trading methodologies have on the market in question as they will differ between asset classes and at different times. The impact of a short or long gamma position in the market can have a big effect on the trading regime of a particular currency pair in the foreign exchange market, for instance, whereas it might not have such an effect on another pair or, indeed, another market. Being aware of the different players in a market and their different trading techniques can really help the understanding of how a market might trade between range trading and trending regimes.

The previous two chapters have examined the role of implied volatility in trading regime analysis but in the next chapter we will look again at an analysis technique that is based on the ever-present humble moving average.

19

My MATE

I love moving averages. They are, like me, simple. They are simple to understand, simple to calculate and simple to work with. However, as I have said, moving averages offer much more than a lagging indicator of trend if the analyst is prepared to be a little creative with them. Take the Moving Average Convergence/Divergence indicator developed by Gerald Appel. The underlying calculation for this very useful indicator is essentially two moving averages of the price and a moving average of the difference between those two moving averages. By looking at the difference between two moving averages the analysis moves from a simple trend (is it up or down?) to the momentum of that trend (is the trend strengthening or weakening?) and this adds an important dimension to the analysis. If we can identify that the trend is up *and* that the trend is strong, it might encourage us to put some more money on the trade; but if we identify that although the trend is up its underlying strength is weakening, it might encourage us to be a little more prudent with a long investment. So, rather than merely using moving averages to identify the *direction* of the trend we can use them to identify the *strength* of the trend, and by doing that it is possible to use a moving average process to identify the current and potential trading regime.

A moving average is a moving average of the price. In most cases it is the closing price that is used but it could be a moving average of the open, high, low, close or any combination of these. It is an average of the previous n periods prices and, as such, represents a smoothed out price series. In this way, moving averages can be thought of as being dynamic trend lines as they follow the trend of the market. The one variable with a moving average is, of course, its length and this will determine how sensitive or insensitive it is to the movements in the underlying price. A 5-week moving average will react far more quickly to the movements in the underlying market price than a 40-week moving average. This is all simple and straightforward but the crucial point is that we would therefore expect that, in a strong trend, a short moving average will be quite a distance away from a longer moving average and, in a weak trend, the moving averages will be much closer to each other.

Taking this logic one step further, we should be able to say that if the difference between the moving averages is greater than some predetermined level, the market could be said to be in a trending trading regime, and if the difference between the moving averages is less than some predetermined level, the market could be said to be in a non-trending or range-trading regime. What we can develop therefore is a method of analysing a market's trading regime based on moving averages, which is similar to the non-continuous trend-following systems like moving average envelopes and the Turtle method. We are able to construct a process whereby we can identify the probabilities that a market is in an uptrending regime, a downtrending regime or a range-trading regime. Since I have been using this methodology in my trading and investment toolbox for years I have referred to it as my MATE (Moving Average Trading Environment).

My MATE consists of a very simple chart of the difference between two moving averages and a cut-off level whereby that difference is deemed to be large enough for the market to be considered to be in a strong trend either upward or downward. When the difference is between the cut-off levels the market is deemed to be in a range-trading regime. Many people use this logic in a slightly different way in trend-following systems. The system rules could be to buy/go long when the short moving average crosses the long moving average from below, but only when the differences between the averages has moved a certain percentage amount away from each other. Conversely, the system rules could be to sell/go short when the short moving average crosses the long moving average from above, but only when the differences between the averages has moved a certain percentage amount away from each other. In this way the system becomes a non-continuous trend-following system because, unless the averages have moved a certain amount away from each other, no trade is taken. Unless the averages move a certain amount away from each other, the market is deemed to be "going nowhere" and is therefore in a range-trading environment.

Some systems use these sort of cut-offs as confidence levels of a trend, so the system might have cut-offs at 1%, 2% and 3% distance that the moving averages are away from each other and, for example, a long position would be taken in one-third of the full position size as and when each of these levels were breached. If the averages are less than 1% away from each other, then there is no position, but as they move up over 1% distance, then a one-third position is taken. When the distance goes over 2%, another one-third position is put on and when the distance goes over 3%, the final one-third of the position is put on. The same logic applies on the way down, so as the averages difference moves back down the long position is gradually decreased. This type of system management is regarded as a money management/position sizing technique and it can be a very good way of improving on a vanilla system where all the risk is put on and taken off at one time. The logic underlying the technique is, however, something that can be used in a multi-fractal setting in the form of trading regime analysis.

to accept the hypothesis that there is a certain amount of persistence in trading regime and that more often than not a trading regime will last for a decent period of time. There are, of course, times when this is clearly not the case. No market analysis is, was or ever will be perfect. However, the times when the trading regime anticipates correctly either an uptrend, downtrend or range more than makes up for the times when it anticipates it incorrectly. What I mean is that, like any good long-term methodology, incorrect anticipation of trading regime (losers) is cut quickly whereas correct anticipation of trading regime (winners) is run.

Due to the fact that it is moving average based, the MATE gives a directional element to the trading regime just as the non-continuous trend-following systems do. This is an obvious enhancement to other trading regime indicators like Bollinger Band Width and ADX where the direction of the potential trend is not as obvious and has to be worked through using other techniques. Where the MATE has an advantage over the Donchian channel-based and moving average envelope-based systems is that it lets the analyst see quite clearly where the points of maximum stretch are. That is, it helps to identify where the difference between the moving averages is very extreme and this can help the analyst in terms of risk management and position sizing. For example, in the EURUSD chart (Figure 19.1) we can see that in October 2004 after the MATE had been signalling a range-trading environment for a few months the market broke higher and the moving average difference went up past the +1% threshold. This movement signals to the trading regime analyst that the range-trading period is probably over and that a new uptrending regime is probably under way.

As things turned out this was indeed the case with the EURUSD moving up in a very strong trend-like fashion from around 1.25 in October 2004 to around 1.36 at the end of December 2004, but at that point in December 2004 the analyst could see that the difference between the two moving averages was at quite a high level relative to the history of the series. This then acts as an indicator that the market could be thought of as being overbought but, crucially, it does not mean that the market has to fall from those levels. It only serves as a guide that the probability of further strong appreciation of the EURUSD is falling. Nevertheless, the fact that the MATE was showing a situation where the positive difference between the moving averages was very high might encourage the trader or investment manager to “skim” some of any long position they had established on the back of the break above the 1% thresholds into the trending regime. Some people use these sorts of indicators to actually sell short or go long as they use them as contrarian indicators, and sometimes they do work like that very well. However, for me the trend should always be assumed to be in force as long as the difference between the moving averages is above the +1% threshold or below the -1% threshold, and using the extremes in the differences to adjust position size, rather than reverse the position, is more comfortable for my particular philosophy (i.e. trend following).

Figure 19.1 shows the euro spot exchange rate with the US dollar or the EURUSD in market jargon using the weekly time fractal. It shows the exchange rate as the solid line and the difference between a 5- and 40-week moving average as the area chart above and below the zero line. This is the MATE (Moving Average Trading Environment) and the reference axis is the left-hand axis that shows the difference between the moving averages expressed as a percentage. There are horizontal lines at 1% and -1% and when the indicator is between these two cut-offs the moving averages are very close together and therefore the market is deemed to be in a range-trading regime. If the indicator is above the +1% level, the short moving average is more than 1% above the longer moving average and an uptrending regime is assumed to be in force. If the indicator is below the -1% level, the short moving average is more than 1% below the longer moving average and a downtrending regime is assumed to be in force.




Figure 19.1 EURUSD MATE weekly

We can see from the chart that when the indicator is in between the cut-offs for a period of time, such as in 2004, then the market does trade in a distinct trading range. This, of course, is just as we would expect given the fact that the moving averages are driven by the price action itself! To reiterate, the crucial aspect to this analysis, and in fact to any analysis, is that we must be able

SUMMARY

This chapter has examined a technique for analysing the trading regime of a market using tolerances for the difference between two moving averages. The logic is similar to that used for moving average envelopes and is similarly flexible.

In the next chapter we will take this trend-following logic a little bit further and consider a trading regime analysis technique that is, once again, simple in its logic but useful in its application.

20

Trend-Following Performance Indicator

“The obvious is always least understood.”

Prince Metternich

This chapter looks at a trading regime indicator that I have labelled the Trend Following Performance Indicator or TFPI. The concept behind the TFPI is very simple and grew out of frustration with searching through ever more complex methods of trading regime analysis. I have always found that the simplest method of analysis provides the greatest clarity, even though it may not be the best all the time, and this idea came to me in one of those simplifying moments. If we break down what we are trying to achieve in trading regime analysis then it boils down to whether the market data suggests if we are in a trending regime or a range-trading regime. Furthermore, we can say that if the market is not in a trending regime, it is more than likely to be in a range-trading regime. Therefore, if we can identify that the market is definitely in, or not in, a trending regime we should be able to identify if the market is in a trending regime or a range-trading regime.

How do we identify if the market is definitely in, or not in, a trending regime? As I said previously, I have always been a big believer in simple facts or methods, which is probably a main reason why I am so attracted to trend following as an investment strategy. As one technical analysis legend is reputed to have said, “Technical analysis does not have to be complicated. Look at the price on a chart. If it is going up that’s good. If it is going down that’s bad.” Statements like this appeal to me because they clarify what is important and don’t waste time with any subjectivity. As alluded to earlier, the market might not always be right but the price is always, always a fact, and if it is going up – even if we are bearish – who are we to argue with it?

With this in mind I became attached to the notion that, when wondering what sort of trading regime the market is in, a simple answer to this question would be to

see if a trend-following system was making money or not. The idea is that if a trend-following system is making money, then the trading regime analyst can conclude that the market is in a trending regime, and if a trend-following system is losing money, then the trading regime analyst can conclude that the market is not in a trending regime, and is therefore probably in a range-trading regime. Fair enough logic but two questions come immediately to mind. The first is “what parameters should we use for the trend-following system?” and the second is “what parameters should we use to decide if the system is making or losing money?”

I decided that the rules of the trend-following system I was going to use should be as simple as possible, so I opted for a very straightforward one rule, continuous system. The trend-following system rules are:

- Open long if the market close is above the moving average.
- Close long if the market close is below the moving average.
- Open short if the market close is below the moving average.
- Close short if the market close is above the moving average.

This is a very simple, straightforward trend-following system but I still had to decide the length of the moving average. When deciding what parameters to use for any market-trading system I think the most important thing is to be logical about it rather than spend a whole lot of time trying to optimise the parameters. My view on optimisation is that it is, despite what the quantitative community would say, essentially data fitting or data mining. Data fitting or data mining are expressions that refer to the practice of trying to “fit” a system to match the data rather than to decide on some fixed parameters and see how the system performs on the market data. Early on in my career I did a lot of optimisation work on various trading systems but I eventually realised that what I was doing was essentially denying the fact that certain trading systems and methods will work in certain trading regimes but not in others.

Optimising, I discovered, was merely a way of trying to make a trading system perform as well as it could in *all* trading regimes and, in my opinion, this is unrealistic. In my view most trend-following systems will make money when the market is in a trending regime and most contrarian systems will make money when the market is in a range-trading regime, so does it really matter to continually optimise the parameters to the best fit over the recent past? By optimising we are making an assumption that the recent *precise* behaviour of the market will continue into the future but I think, while it is realistic and logical to assume that the general behaviour of the market (e.g. trending or range trading) will continue into the future, it is not logical to assume that the absolutely *precise* nature of this behaviour will continue into the future. However, this is what the optimisers are doing. There is also the question when you optimise system parameters of how often you re-optimize them. Is it every week? Every month? Every quarter? The decision becomes subjective, so what’s the point of optimising in the first place? And of course the biggest question of all is how

much risk to put on the system or the individual components of the system? Equally weighted, subjectively weighted or optimised weights? The issues are the same.

I decided therefore, when thinking about what parameters to use for the Trend Following Performance Indicator, that I would use a 21-period moving average because that period was neither too short to miss the long-term moves nor too long to miss the short-term moves. It could have been anywhere between 15 and 30 but I decided on 21. I realised that 21 was a Fibonacci number, and while I do not believe that there is some “magic” in Fibonacci numbers, I do believe there is a logic to the subject of natural mathematics that the Fibonacci number series represents, and for that reason I have an emotional attachment to them (see Chapter 8 for an explanation of the Fibonacci number series). The precise number of the moving average is not important; just that it is neither too short nor too long a moving average.

The trend-following system is therefore to go long when the closing market price is above the 21-period moving average and go short when the closing price is below the 21-period moving average. This is a very simple system but is the purest essence of a trend-following system because, being a continuous system, it always will have a position in the market either long or short whereas if it was a discontinuous system it would have an element of anticipating a range-trading regime. Always having a position in the market means that it is always following the 21-period moving average. If there is a trend it will make money and if there is not a trend the result will be a lot of “whipsaw” and it will lose money. For the purposes of the TFPI this is the behaviour of the system that I wanted.

The next question to address was how to decide when the system was making money and when it was losing money. As we know, any system will go through periods of winning and periods of losing (or drawdown) but there are various ways to look at statistics when deciding if the system is currently winning or losing. Should we look at returns year to date, month to date or since inception? Should we look at whether the equity curve (the profit and loss curve) is above or below the most recent peak? The problem with these measures is that they tend to be fairly static. What I wanted was a dynamic measure of the profitability of the trend-following system and it was for this reason that I decided to look at a moving average of the equity curve.

The equity curve, or the profit (and loss) curve, for the system is just the time series of the cumulative profit, but by relating it to a moving average we can get a sense of whether the system is “working” (making money) or “not working” (losing money or experiencing drawdown). If the equity curve is above the moving average then we can say that the trend-following system is making money and if the equity curve is below the moving average then we can say that the trend-following system is losing money. Furthermore, if we can say that the trend-following system is making money, we can say that the market must be in a trending regime, and if the trend-following system is losing money, we can say that the market must not be in a trending regime and is therefore, in all probability, in a range-trading regime. This, then, is the logic of the TFPI.

This logic of looking at dynamic performance using moving averages of the profit and loss or equity curve is not really new. It is actually how a lot of proprietary trading desks allocate risk budget (the amount of money a trader is allowed to trade with). Let's take an example of an investment bank that has a "prop" desk with three traders whose mandate is to trade the currency markets, and the maximum allocation is risk of, for instance, a US dollar amount of 90m. Each trader starts off with a risk budget of USD 30m and let's say trader A makes a lot of money over the first few weeks, trader B's performance is fairly flat while trader C is making some losses. The risk budget allocators would look at this and maybe do nothing. They might say, "let's wait until the end of the year to see how performance looks then and re-allocate the risk budget at that time". A more dynamic system, however, might be to allocate more now to trader A because he is making money and allocate less to trader C because he is losing money. In this regard the risk is being increased to someone who is outperforming and decreased from someone who is underperforming. It is an anti-martingale method of money management whereby risk is increased in line with the growth of the equity curve.

The martingale system of money management is based on the betting premise that each time you lose you should double the bet size because eventually you will win and recoup all your losses. While this is mathematically true it is unrealistic in the real world because it assumes that you have infinite resources at your disposal. Many people have found out the hard way that using this money management strategy can be ruinous when markets have extreme trending moves. In the words of the statisticians, this is when the market goes to the tails of the (normal) distribution, when the rare events occur. You see, the martingale strategy of money management relies upon the market staying within the fat part of the normal distribution (i.e. that it will mostly mean revert and not have any large, extended trends). As alluded to earlier, the move to the tails happens more often than we think, and when it does it is the trend followers who make money and the mean reverters (those who expect the market to mean revert) who get clobbered.

An anti-martingale money management method, however, takes the opposite view and increases the bet size as things are going well. The philosophy here is that you will be adding risk when things are going well so that should increase your equity exponentially. This, of course, means that when you get the inevitable loss it will occur when you have your biggest bet on (at least in nominal terms).

The difference between the two methods is the assumed risk/reward ratio. In the anti-martingale method you have to assume that when you win you will win more than when you lose in order for the method to work. If someone came up to me and asked me to play a game of coin toss and when I won I would win £1 and when I lost I would lose £1, would I play this game? Well, I could, but it would be a complete waste of time because over the long term I would end up with a flat account. That is zero. The long-term probability of a game of coin toss is 50% – 50% of the time it will be heads and 50% of the time it will be tails – over the long term it has to be stressed. Now, what if the same person asked me to play the game of

coin toss but this time when I won I gained £2 and when I lost I lost £1 like before? Would I play this game? Let me tell you, if anyone asked me to play this game for the rest of my life I would bite his hand off to do it because this is a guaranteed way to build a monetary fortune! In this game 50% of the time I make £2 while the other 50% of the time I lose £1. Over one million tosses I could assume that my account would be up £500,000.

In a sense, however, it pays to take account of both a martingale and anti-martingale method of money management because at some point the outperforming trader or method will start to underperform and vice versa. The performance of the trend-following system will reach extremes on both sides of the profit and loss line, and when it does it would pay to either increase an allocation to trend following in the case of underperformance, or decrease an allocation to trend following in the case of outperformance. Identifying these peaks in performance is part of what the TFPI seeks to do, and it is at these extremes that a potential shift away from one regime to the other should start to appear.

For trend followers one of the facts of market life is that there will be periods when the strategy underperforms because markets are just not trending strongly. When you accept trend following as a philosophy you have to accept this otherwise you will tie yourself up in knots of frustration wondering why the strategy does not perform in each and every period. However, the interesting aspect about this is that *trend following is a long-term strategy* even though most people would probably lump in trend following with a perception of short-term technical trading. What I mean is that, although there may be periods of underperformance as markets do not trend for a while, over the long term the strategy will make money simply because it is the best and most consistent strategy for adhering to the age-old investment philosophy of running winners and cutting losers. Nevertheless many people forget about this when they see performance of trend following, particularly when trend following is underperforming and losing money. Most trend followers will have stories of all the criticism and comments they get when going through periods of drawdown to the effect that “markets have changed” and that trends no longer exist.

In a very real way this psychology of dismissing trend following after it has been going through a period of underperformance, and indeed the opposite reaction of overhyping trend following as it has been going through a period of outperformance, is just exactly the same as any psychology cycle for a market going through bull and bear cycles. At the top, the psychology is at the zenith of bullishness and, at the bottom, psychology is at the nadir of bearishness. It makes me laugh when people say that markets have changed and that trend following is dead. They are saying that trends will not be the same as they were in the past and a lot of our quant friends tell us the “reason” for this is that the strategy has been “arbitraged away”. What they mean is that because so many people are using trend-following strategies, the “alpha” that was present is no longer available as the market has become too crowded with people trying to do the same thing. The problem with this argument is that it

assumes that most trend followers will be following the same or similar systems in one particular time fractal and, while I can accept that there are a lot more trend followers now than there used to be, they do not all use the same system or, more importantly, money management methodology. The other problem with the idea of trend following being “arbitraged away” is that it assumes that trends are something that is an arbitrage opportunity in the first place. An arbitrage opportunity exists when close or similar financial instruments or markets are displaying a different price that enables market players to come in and exploit this difference by buying one and selling the other. The more money that does this drives the prices back in line with each other and, in that case, the opportunity can be said to have been arbitrated away. However, I do not see how the same hypothesis can be used for a market strategy that relies on market behaviour that is based on human psychology such as trends.

One of the things I love about the financial markets and investment is that the old adages we hear about turn out to be so true. Adages like “run your winners, cut your losers” and “the first cut is the cheapest” have stood, and will continue to stand, the test of time. This is why it pays, in my opinion, to listen to the older market practitioners. People who have been around market cycles over the years have the best understanding of what makes markets tick, how certain things change and how certain things do not change. People like Richard Russell of *Dow Theory Letters* are the people that we should all listen to. He started writing his market letter in the 1950s and, at the time of writing, in 2008, was still going strong. The words of wisdom about market behaviour that come from people like Russell are, in my opinion, invaluable to anyone involved in wanting to trade or invest successfully in the markets. Over the course of his decades in the business, Russell has realised that, as he says, “everything works and nothing works”, meaning that the ebb and flow of the markets will enable all methodologies to go through some periods of out-performance only to underperform in other periods. My take on what he is saying is that in order to be successful, one should follow a methodology that, although it will go through periods of underperformance, possesses an ethos and philosophy that captures the *long-term moves in the markets* that will happen eventually. In his case this philosophy is Dow Theory which, in essence, is a trend-following methodology albeit mainly discretionary rather than mechanical. Perhaps the single most poignant quote about what changes and what does not is attributable to Jesse Livermore, the legendary trader who experienced many different cycles in his time, when he said: “Wall Street never changes, the pockets change, the suckers change, the stocks change, but Wall Street never changes, because human nature never changes.”

Long-term trends happen. They happen because human nature drives prices, and they happen at the same time as fundamental and economic factors are changing. The point of being a trend follower is not to deny that long-term trends and shifts in market prices will occur at the same time as fundamental and economic factors are changing, but just that the reasons for those changing economic and fundamental factors only becomes known *after* the market price has already moved! Perhaps

trend followers should call themselves economic followers or fundamental followers because that is essentially what the strategy is doing. The price moves coincidentally with, if not before, the economic and fundamental factors change and so following the trend of, for example, a currency is exactly the same as following the changing economic fortunes for that currency. People who say that trend following is dead and that trends will no longer happen are actually getting into a debate about economics and politics. If someone thinks that the world has become one state where no change happens in the realm of economics, politics and social mood, then fine, I am willing to concede that it would be hard for markets to trend. However, I am willing to bet that as long as human beings exist, social mood will ebb and flow, wars will be fought and peace will be made, economies will go through boom and bust and change will constantly occur. These things mean trends in markets.

It is ironic, therefore, that there is a mean-reverting nature of performance when it comes to a strategy like trend following, and this is simply reflecting the observation of volatility regimes that markets will trend and then range before trending again. What this means is that the best time to invest in trend-following funds, or the best time to allocate more risk to a trend-following system, is when the performance of that trend-following fund or system has been quite negative. In this sense, it could be said that a martingale type of philosophy should be followed whereby you expect performance to mean revert. Importantly, however, these mean reversions will happen at the extremes of performance.

Figure 20.1 shows the Trend Following Performance Indicator (TFPI) for the EURUSD exchange rate between 2002 and 2008. The right-hand scale shows the TFPI; when the shaded area is below zero the trend-following system's equity curve is below its moving average (a range-trading regime) and when the shaded area is above zero the equity curve is above its moving average (a trending regime). This TFPI is using a 21-week moving average for the trend-following system, a 10-week moving average of the equity curve and a 20-week moving average of the difference between the equity curve and its moving average. The TFPI shaded area is the actual difference relative to its moving average. There are two things to notice. Firstly, the change between regimes can be very quick and can sometimes occur with a lag. Secondly, however, the extreme levels of regime can on many occasions preempt the switch from one regime to another. In October 2004, for example, the TFPI turned positive indicating that the equity curve trend of the trend-following system was moving higher and this pre-empted a trending move in the EURUSD exchange rate from 1.24 to 1.36 by December of that year. At the start of December the TFPI had blasted through the previous high levels that had been achieved, indicating that the momentum of the trend-following performance was beyond the extremities of the history and, at this point, the trading regime analyst would have been cognisant of the fact that the chances of the current trend maintaining its momentum were decreasing rapidly.

In this case the TFPI can be used in a mean-reverting fashion in that the analyst would be thinking it would be prudent to at least lock in some gains on the trend now




Figure 20.1 Trend Following Performance Indicator (EURUSD weekly)

that the market appeared primed to start the regime-switching process. However, like any mean-reverting type of indicator that exists to indicate an overbought or oversold condition like relative strength or stochastics, the sensible thing to do is to wait until the indicator stops going up, even though it is in record extreme territory already. The one thing you do *not* do in the markets is the financial equivalent of standing in front of a speeding train that is still accelerating! We can see that once this record momentum of the equity curve started to abate, the TFPI moved back down below the zero line at the start of January 2005, indicating that the equity curve trend of the trend-following system was moving lower and this pre-empted a choppy, range-trading period until early May when the exchange rate moved in a range of between 1.27 and 1.35. In early May the indicator turned up again, indicating that the trend-following performance was gaining momentum and this time the market moved down in a strong trend from 1.27 to 1.1850 by July.

As usual, the question arises as to what is an extreme level of the TFPI. I decided not to put any upper or lower boundaries on this indicator but, of course, that is easily done. For this indicator I decided to just let the upper and lower bounds fall naturally, so there will be times, such as in December 2004, that the TFPI blasts through the previous extremes. However, in my opinion, it is best to wait for all oscillating type of measures to stop going up before acting in any contrarian manner. If we do this, we can avoid the analogy of standing in front of the speeding train.

SUMMARY

The Trend Following Performance Indicator is not perfect – far from it. There will never be a perfect indicator of anything. However, the TFPI can give the trading regime analyst a very good idea of where the probabilities lie in terms of the likely trading environment for a market. Not only does the TFPI give a measure of the trend of the trading regime, and therefore its longevity, it also gives a measure of how extreme the trading regime has become and therefore an indicator of a possible regime reversal.

In the next chapter we will examine an indicator that takes us full circle back to the use of the standard deviation for trading regime analysis.

21

Trading Regime Indicator

When I first started thinking about and examining trading regimes in terms of trending or not trending (ranging), it was quite obvious that it was reasonably straightforward to identify where a trend was likely to start by looking at analyses like orthodox pattern recognition and identifying breakout points. For example when one identifies that the market in question is tracing out an ascending triangular pattern, then when the upside break does come it is implicit in the analysis that there will be a switch from the ranging regime just seen to more of an impulsive or trending regime. After range contraction comes range expansion. Equally, it became evident that identifying the likely end of trending periods could be done by analysing volatility extremes such as with the Bollinger Band Width. While perhaps not always resulting in a period of distinct range trading, volatility extremities do a good job at identifying trend exhaustion points.

It appeared, therefore, that the tools to identify the potential start of trending regimes, and to identify when those trends might have reached an extreme point, were available to the trading regime analyst already. What I thought would be useful, however, would be a method to identify not only the start and the extremity of trending periods but something that helps to identify the longevity of those periods and, by default, the longevity of range-trading periods also. Non-continuous trend-following systems like the Turtles used do an excellent job in this regard and in order to identify the longevity of trading regimes we can also look at the *direction* of the Bollinger Band Width or the ADX to determine whether the market range is contracting or expanding. If the indicator of volatility is going up, the logic suggests that it is more likely to be a trending regime, and if the indicator of volatility is going down, it is more likely that the market is vulnerable to a range-trading regime. However, this was quite a subjective exercise and what I wanted to see was a more objective measure of the trend in volatility itself.

The Trading Regime Indicator (TRI) is a simple indicator that I have developed which combines a volatility measure and an objective trend of that volatility measure.

Ideally what we are looking for in a trading regime indicator is something that identifies when there is extreme regime behaviour but that also identifies the trend of that regime behaviour in order to extract as much value as possible from the analysis.

The TRI is calculated by taking the standard deviation of a price series, taking a moving average of that standard deviation, and examining the gap between the two.

In some senses the TRI is quite similar to the Bollinger Band Width but instead of looking at the width of standard deviation bands away from a moving average the TRI is looking simply at the standard deviation of the closing price itself. Furthermore, whereas in the sections on Bollinger Band Width and the ADX, we examined the usefulness of looking at the trend of the band width and the ADX from a discretionary perspective, being objective about the definition of the trend in the standard deviation is something I have found to be quite useful. It does, however, present its own challenges when applying moving average analysis on an underlying series that itself is quite volatile in the sense of moving up and down a lot over short time spans.

Figure 21.1 shows the line chart of the weekly closing levels in the US Treasury 10-Year Bond Yield between 2002 and 2008. Overlaid on the chart is the Trading Regime Indicator, which is the shaded area fluctuating above and below the zero line on the right-hand scale. This TRI is using a 20-period standard deviation and a 200-period moving average of that standard deviation in the calculations, and the shaded area shows the difference between the standard deviation and the moving average expressed as a percentage of the average. When the TRI is above the zero line




Figure 21.1 Trading Regime Indicator – US Treasury 10-year Yield weekly

it indicates that the standard deviation is above its moving average and that a trending regime is in existence. On the other hand, when the TRI is below the zero line it indicates that the standard deviation is below its moving average and that a range-trading regime is in existence. The horizontal lines are set at the +50% and -50% levels to indicate where the regime could be considered to be in an extreme mode (i.e. that there is an extremely strong trend or that there is an extremely persistent range).

We can see from the chart that, in general, the TRI does a good job in indicating the current trading regime but, as always with using moving average based indicators, there are times when the lag involved in switching above and below the zero line hurts the analysis. This is simply a fact of life with something that is attempting to identify a trend; therefore, when we are using a trend-based analysis such as this, the overriding belief that we must have is that whatever it is we are analysing will exhibit some degree of persistence over time. What I mean is that if we are using a trend-based measure to gauge the trading regime, then we have to believe that when a regime exists it can go through periods of persisting in that regime for some time. This is exactly the same philosophy of trend following in general whereby we fully expect to have times when the trend signal turns out to be a false dawn, but that when it isn't a false dawn the longevity and extent of the move allows us to more than make up for the losing times. The winners are run and the losers are cut.

We can see that in March 2002, towards the left-hand side of the chart, the 10-year bond yield was experiencing a bit of a choppy period but then, when the market started to trend down the TRI switched from a range-trading mode to a trending mode. It did so on the week of 12 July 2002 when the yield was trading at 4.57% and the market subsequently fell in a strong trend-like fashion down to a low of 3.66% on 4 October 2002. That is almost a full percentage point drop in yield in four months, which is equivalent to over a six-point move in the bond futures contract. Clearly, although the lag of the indicator meant that it missed a little bit of the start of the trending regime, the TRI enabled the trading regime analyst to identify this as a trending regime in enough time to capture the meat of the move. And this is what trend following does. The philosophy assumes that you are not going to capture the exact start or the exact end of a move in price or, in this case, a switch in the trading regime, but that what you capture will be enough to add value to your process. No person or mechanical system will ever capture the exact start and/or end of moves in the markets all the time, so in my opinion it is not worth wasting your time trying. We might as well use something that we know will guarantee us to be on the correct analytical side of a persistent regime. Again, all we can do is talk about probabilities, and in this case the probabilities were that the trading regime had switched to a trending regime, so if we are trading the market we should be cognisant of that fact. As this is a weekly time fractal chart, we could take this information and use it when we actually trade on the shorter time fractals. That is, if we were trading using daily or hourly charts, we should be factoring in that a higher fractal trending regime is in force and should, perhaps, be looking to go with breakouts rather than trying to fade them.

Although no person or mechanical system will ever capture the exact tops and bottoms all the time, what we can do is use extremity observations to get a sense of where exhaustion points might occur. This is not trying to call tops and bottoms, but is merely more information we can use to factor into our decision-making process. For instance, by the end of September 2002 in the example of the US 10-year yield (see Figure 21.1), the TRI had reached the 80% levels, indicating to the analyst that the distance of the standard deviation away from its moving average was very extreme. However, rather than say “that must be the turning point and the beginning of the end of the trending regime” the analyst is best to use this information in a way that tones down their trending bias rather than switches them, then and there, to a range-trading regime bias. This is consistent with the age-old reasoning that just because something is overbought or oversold it doesn’t mean that there will be a market reversal. It increases the odds of it, but it doesn’t mean there will be a reversal because the market can go sideways in order to “work off” the high oscillator readings. As it turned out, the extremely high readings of the TRI at the end of September 2002 did turn out to be the start of a range-trading regime, and that regime switch was confirmed when, on 22 November 2002, the TRI moved down below the zero line. That presaged a choppy, range-trading period for the next few months when the market price action provided many opportunities for a mean-reverting trading style to make money.

Like all regime indicators, the Trading Regime Indicator can also be used for shorter time fractal analysis, including intra-day. However, if you are a trader/investor, or if you are putting recommendations on if you are an adviser, it can be a very useful application of trading regime analysis to use it for at least one time fractal above the fractal that you are actually trading the market. That is, if our investment horizon is longer than a few weeks, it would be best to use an analysis of volatility based on the monthly time fractal to come to a conclusion about trading regime, but if our investment horizon was shorter than that, we could use a volatility analysis using the weekly time fractal. What about very short-term investment horizons? Well, for people who like to trade the markets on a shorter-term basis (perhaps using hourly charts to trade with), the use of a volatility analysis based on the daily time fractal might be best.

Figure 21.2 shows the Commodities Research Bureau (CRB) Index using the daily time fractal from April 2006 until July 2007 with the 20/200 Trading Regime Indicator overlaid. The 20 refers to the 20-period standard deviation and the 200 refers to the moving average of that standard deviation. We can see from this chart that the TRI can get whipped around a little, as it did in the middle part of 2006, having identified a strong trending regime before that in April. For short-term traders using this analysis, it would have been a frustrating time as the indicator went from trending to ranging and back again quite a few times. The reason for this is that the market itself was trading in quite a large range and this range was just too large for the daily time fractal TRI to take account of it. This is unfortunate, but I want to get the honest point across that the TRI, as is the case with all indicators, is not infallible and it will




Figure 21.2 Trading Regime Indicator – CRB Index daily

go through periods of giving false signals. I think most honest analysts will testify to this fact but as it sometimes gets lost in the analysis, we must stress time and again that we are dealing with probabilities and all any indicator does is give a probability reading. There are *no* certainties in the markets, ever.

However, once this frustrating period had passed the CRB Index went back to more traditionally persistent regime and we can see from the chart that, had a short-term trader been using this analysis to decide position size and/or whether to stick with a trending breakout or to fade it, the indicator did a good job in helping the analysis.

SUMMARY

In this chapter we have looked at the Trading Regime Indicator, which is a fusion of standard deviation and moving average analysis. Like all indicators, it is far from perfect but it does, I think, provide a logical and consistent basis for evaluating the volatility conditions of a market and can give the analyst both an idea of the current trading regime and, when that regime is at extreme points, a potential turn into a different trading regime.

Part IV

Combining and Using Trading Regime Analysis

“The more I practice, the luckier I get”

Gary Player

22

An Eclectic Approach

So far we have discussed how a market for a stock, bond, currency or commodity is driven by human psychology and how that psychology can swing in cycles that will affect the trading regime (trending or range trading) of the market. We have examined the existing indicators that can be used to help us to identify the current and probable future trading regimes, as well as touch upon some further ideas that can assist us. However, to use another golfing analogy, it's nice to have well-made and personally fitted clubs in your bag, but if you don't know how to play the game the clubs are pretty useless. That is, it is all very well for us as trading regime analysts to have the appropriate indicators and methodologies for identifying trading regimes, but if we don't know how to combine and use them, we won't be realising their full potential for helping returns in the markets. This section deals with the practicalities of using trading regime analysis and examines situations whereby combining individual analyses can increase the probabilities of success in actually trading the markets.

The basis of this book is that not all trading or investing methodologies will work in every single time period and so, by analysing the current or likely market conditions, we can better trade or invest the markets. Being pragmatic and realising that there a number of ways to achieve goals, and that not one mantra is appropriate on every occasion, is a realistic way of getting through not just the markets, but life itself. Market analysis is like solving a puzzle. The more pieces of evidence that point to a certain outcome, the more probable it is that this outcome will be the most accurate. It's all about probabilities. In fact, for most of us the human brain acts naturally as a sort of probability machine when it is asked about market direction. If someone is asked "what do you think of the outlook for the stock market?" his iterative neurological process would be to assimilate the factors that might affect the stock market, assess the current as well as potential future conditions and attach weights or probabilities to each factor. This does not mean that people are immune to the psychological flaws discussed earlier, such as cognitive dissonance, but just that for most people this is how they would come to a decision about the market in the first

place. In one sense, we can think of this natural thought process as a very loose form of Bayesian probability logic, named after the great mathematician Thomas Bayes (1702–1761), whereby weights and probabilities are attached to various outcomes in order to come up with an overall probability.

Using this logic in terms of trading regime analysis can be very useful indeed because, as we have seen, there is not one trading regime indicator or method that will give a guarantee of success every time. However, by *combining* our analyses together we can create an overall probability reading of success or failure in identifying probable trading regimes. The more mathematically minded of us could create an objective process for doing this or we could do this via a simpler, discretionary process of having a table showing each trading regime indicator or method in one column, and a column to describe the particular regime that the time fractal is favouring. An example of such a regime is shown in Figure 22.1, which gives an example of what a trading regime analysis might look like using this Bayesian type of logic on the US Dollar spot Index. Although the US Dollar spot Index, based on the US Dollar Index future (DXY) has quite a high weighting to European currencies versus the US dollar, it can be used as a good proxy for the overall trading conditions of the G10 (Group of Ten) currency markets. That is, by analysing the US Dollar Index it enables us to anticipate the likely trading regime for the G10 currencies overall.

We are going to examine the period from 1999 until 2003 and our starting point is October 2000. The Dollar Index had been in a long rally from 1998 when it was trading around 90 and at October 2000 stood around 118. Clearly it had experienced a solid trending period but was the trading regime analysis signalling that the trending period was likely to make way for more of a range-trading regime?

If we go through each of our regime criteria we can get an overall idea of the probabilities involved, so starting with orthodox pattern recognition, there was, in October 2000, no discernible pattern that would make us think that the trending period was about to end although the Elliott Wave counts were both indicating mature fifth waves. The monthly (structural) candlestick for the month of September 2000 was a “spinning top”, meaning that the open and close were relatively close together around the middle of the high to low range, which was itself relatively large. This is a signal of uncertainty in the market and coming after a trend it signals that psychology is shifting away from the powers behind the trend. In terms of previous highs and lows it would have been of interest to the analyst that the 120 zone was where the Dollar Index effectively started from at the break of the Bretton Woods system of currencies in the early 1970s. In other words, the 120 zone was a natural resting point for an uptrend that lasted eight years from the lows in 1992 around the 78 level.

The 10-period Rate of Change in the structural (monthly) time fractal was towards the top end of its historical extremes, as was the same indicator using in the cyclical (weekly) time fractal. The 20-period Bollinger Band Width using the structural fractal was at previous highs, seen in both 1998 and 1995, and the cyclical fractal

	Weight	Structural	Cyclical
Orthodox Pattern Recognition			
Elliott Wave Analysis		R	R
Bollinger Band Width		R	R
Moving Average Envelopes			
MATE		UT	UT
Donchian Channels		UT	UT
ADX		R	R
Volume			
Previous Highs / Lows		R	
Williams % R		UT	UT
TRI		R	
TFPI		R	
Volatility Smile Index			
Nerve Index		R	R
ATR		R	R
Rate Of Change		R	
Candlestick Analysis		R	
Point & Figure			
Conclusion		R	R

Figure 22.1 DXY Trading Regime Grid – October 2000

volume considerations, gap theory and most other methodologies will all be cousins of each other. However, in my opinion, the linkage through which this takes place – in the final analysis the purest form of price action analysis – is the Elliott Wave Principle because the Elliott cycle describes the underlying driver of all price action. Therefore, using Elliott as a pillar of your trading regime analysis is, in many respects, the sensible thing to do.

All these linkages are inevitable because we are, after all, analysing the same underlying variable, the price; it is just that there are different ways of expressing the analysis. For me, it all boils down to volatility. Either price ranges will be expanding or contracting, but the clues to analyse the potential phase for volatility come from many different aspects of technical analysis. Just as a police detective would go about his business, by doing the hard work of covering all aspects of the case and searching for clues, at some point we can come up with a confluence of evidence that gives an overwhelming bias towards a criminal charge being brought.

Some more examples of combining analysis techniques follow and, although by no means exhaustive, they can give the reader a feel for how certain aspects come together to create higher probability scenarios.

In February 2008, the weekly time fractal in EURUSD was showing low Bollinger Band Width and an Elliott corrective fourth wave that had started in November 2007. Figure 22.3 shows the weekly bar chart of EURUSD and we can see a clear and valid Elliott wave count with wave iv unfolding in what also appears to be an ascending triangle.

This volatility or range compression is the set up for a volatility breakout trade using a lower time fractal. We could add to the analysis that, although not shown, the structural (monthly) and cyclical (weekly) Supply Demand Indices (moving average analysis, see Appendix 1) was positive, that the daily ADX had been falling and was extremely low and that the price had drifted back into the 1% moving average envelopes. All of this evidence would have suggested to the analyst that the trading regime was clearly ranging, but furthermore that the regime was reaching some sort of extreme. Low volatility begets higher volatility, but which direction does the price go with that volatility? In this example, given the wave count, the ascending triangle and the fact that the supply and demand mix was positive (short moving averages were above longer moving averages) the analyst could skew the probabilities in favour of the resolution into higher volatility being to the upside in price.

Therefore, when it came to actually trying to make some money out of the analysis and putting on a trade, we would drop down a time fractal and use the daily data for our trade risk parameters. This is because when we are trying to take advantage of a volatility breakout such as this we are, as the name suggests, buying volatility. We are anticipating a strong trend in the price, but given that it can be extremely difficult to anticipate where a strong trend in motion will stop, it makes sense to play the volatility breakout by utilising a trend-following methodology and use that as a method of risk management. We could use a simple daily moving average system to enter the trade and also as a trailing stop-loss that would keep us in the trend while,

Bollinger Band Width was also at previous high levels. The Williams %R readings were both showing that an uptrend was in full force and had not yet turned back into the neutral zone. In October 2000 the monthly 14-period ADX was at a very high level of 37 while the weekly ADX was above 40. The Average True Range was at 12-year highs on the structural fractal while the cyclical ATR was also showing multi-year highs.

Put together, all of these analyses amounted to an overwhelming amount of evidence that the period of the trending regime was extremely mature and that in all probabilities the next period would be one of range trading. As we can see from Figure 22.2, the US Dollar Index subsequently traded in a tight range from around 120 to 110 over the next year and a half, giving currency traders or currency overlay managers the opportunity to take advantage of these range-trading conditions by either trading in a mean-reverting manner or by allocating more risk to short volatility strategies like option selling and less risk to long volatility strategies like trend following. During this period trend-following vastly underperformed those strategies like option selling but also including carry trading, which take advantage of range-trading regimes.




Figure 22.2 US Dollar Index monthly

Most technical and volatility analysis techniques and methods will be linked in some way. Just as low volatility will probably be associated with a coiling price pattern such as a triangle and high volatility will probably be associated with a trend exhaustion turning phase perhaps with momentum divergence, candlestick reversals,




Figure 22.3 EURUSD weekly and Bolling Band Width

at the same time, enable us to take our profit before a reversal in the trend hurts us. In this sense, the potential trade holding period could be multi-week, depending on the strength of the subsequent volatility breakout.

Figure 22.4 shows the daily time fractal in the EURUSD exchange rate and how the volatility breakout unfolded. By the middle of February 2008 we would have been primed for a volatility breakout probably to the top side and so when the 5-day moving average crossed above the 40-day moving average on 20 February we would have initiated the trade plan by putting on some of the position (say 50% of our total risk allocated to the trade) at 1.4715. The risk management of that position could have been to stop-loss out of half of it if the moving averages turned back down, and stop the rest of it below the triangle at 1.4450. Given the fact that we are using the moving averages as a trend-following risk management system for this trade, then we must be prepared to continue to use them as such until it is clear that a trend has taken place. That is, we are identifying a zone where there has already been some




Figure 22.4 EURUSD daily

consolidation and choppy price action, thereby making the assumption that chop will be limited from here on and a trend will develop. The use of a disciplined trend-following methodology to actually take advantage of that means that we must be prepared to trust that methodology for some time. If the averages turned back down and stopped us out of half the position, but then turned back up again, we must be prepared to re-enter the trade if the underlying trading regime analysis has not changed. In this case, however, the averages did not turn back down and on 26 February we would have added to the long position around 1.4970 as the EURUSD exchange rate broke above the ascending triangle. Once the entire position is established, the stop-loss for that entire position becomes the trailing moving averages. The profit target is determined by the trailing moving averages but we could also be prepared to take some of the long position off if volatility became very high. For instance, a high daily fractal Bollinger Band Width on 7 March might have encouraged us to take some profit around the 1.54 level.

In AUDCHF, Figure 22.5 shows that the daily time fractal Bollinger Band Width was quite low by 14 April after a period of consolidation. The Supply Demand Index (not shown here), although negative, was on an improving path from the extreme point.

Moving down the time fractals, Figure 22.6 shows a clear Elliott cycle. It shows the hourly chart of AUDCHF from 31 March until 21 April 2008, and so a potential trading plan using the hourly time fractal (a multi-day outlook) on 15 April would have been to; buy AUDCHF on the moving average cross at 0.9291 (stop 50% of that




Figure 22.5 AUDCHF daily and Bollinger Band Width




Figure 22.6 AUDCHF hourly 31 March–21 April 2008

position on a moving average break back down and the rest below the wave C low at 0.9175); and buy the moving average stop back if they cross back up (which they did at 0.9280 later that day). The stop is the trailing moving averages. The profit target is determined by the trailing hourly moving averages, but when volatility gets very

high (high Bollinger Band Width) then a partial reduction of the position would be appropriate at that juncture.




Figure 22.7 USDCAD weekly and Bollinger Band Width

Figure 22.7 shows the weekly time fractal chart of USDCAD and we can observe that, in July 2007, the exchange rate was experiencing extremely high weekly Bollinger Band Width at the end of a valid wave 3 down (along with extremely high moving average differences which are not shown). We can see that the market then consolidated in a distinct ranging regime for a number of weeks. Then, in November, after a valid fifth wave count the market experienced a distinct hammer bottom candle at multi-year trend line support. On both of these occasions we could have taking advantage of this trend exhaustion by, for instance, establishing short-term USDCAD bull call or put spreads that would benefit from a lowering of volatility and a sideways to up movement in the market. We could also have moved down the time fractals and used mean-reverting techniques such as slow stochastics to benefit from the sideways movement.

In SGDJPY (Singapore dollars priced in Japanese yen), Figure 22.8 shows high weekly volatility (Bollinger Band Width) on the week of 17 March at the end of what could be counted as an Elliott Wave corrective double 3. This is a set up for potential exhaustion of a move.

Looking down the time fractals to the hourly chart (Figure 22.9), we notice that the most recent sharp move down in SGDJPY has done a clear Elliott five waves, so this could signal the final C wave of the double 3.

23

Applications for Traders and Investors

Trading regime analysis has many uses for traders and investors from the very short term to the very long term. There is, of course, a very broad spectrum of market participants and each will have his own particular methodology, time frame, philosophy, money management, markets to trade or invest in and numerous other variables, but the one common glue that unites us all is the requirement for prices to move or not to move in order to make returns from the market. That is, we might all be using a vast array of methods and philosophies in order to predict or anticipate the future price action of the markets, but it is the trading regime of the market that will ultimately determine if our selected strategy is successful or not.

Now here's a thought. What is the difference between a trader and an investor? Let's make it easier. What is the difference between a trader who is a price taker (not a market maker) and an investor? Is a trader someone who deals in the very short term, whereas an investor deals in longer time periods? Maybe. In my book, however, there is no difference whatsoever between the terms because each is doing the same thing. Each person is buying and selling assets or derivatives in order to manage risk. Someone might be at home trading currencies on an internet account and someone might be working for a massive insurance company managing a portfolio of stocks, but there is no difference in what the pair are trying to achieve. They are both trying to maximise return. Fair enough, the portfolio manager might be trying to maximise return in relation to the amount of risk taken (the famous Sharpe Ratio or Information Ratio, which is the darling of institutional investors) but both are seeking to buy and sell positions to create value no matter what benchmark that is against. So why do people differentiate between "traders" and "investors"? Having worked in what you would call traditional, institutional investment companies there always seemed to be a certain snobbery of the fact that they, as investors, are somehow much more professional and serious than so-called traders, who are either




Figure 22.8 SGDJPY weekly and Bollinger Band Width



Figure 22.9 SGDJPY hourly, 10–25 March 2008

On 17 March, therefore, the analysis would suggest a large potential for trend exhaustion setting in, and by establishing short-term bull call or put spreads we could take advantage of the subsequent sideways to up move.

working for themselves as independent traders or managing small hedge funds. A lot of this differentiation is due to perceptions regarding investment time frame of course, where an investor is investing for the long term and a trader is trading for the short term. Whatever the semantics involved, the point is that trading regime or volatility conditions affect both investors and traders alike because of the fractalised nature of markets. Therefore, the applications of trading regime analysis will be spread across all time frames, from the very short day-trading fractals to the very long multi-year investment fractals. Let us look at a few areas where trading regime analysis can be applied.

CHOICE OF MARKET TO TRADE

Self-awareness is, in my opinion, a very good trait to have. If someone is displaying self-awareness in the psychological sense, then that person has the realisation of oneself as an entity and of one's own individuality. People who are self-aware are generally contented people who "know themselves" in the sense of knowing what they like, what they don't like and are confident enough to tread their own path through life. Being self-aware when you are trading the markets is quite essential because the seemingly endless list of ways to extract money from the markets can get people turning back and forth from one method to another in search of the non-existent holy grail. However, a self-aware person will know what his preferred investment or trading methodology is, and when that can be nailed down he can treat any other methodology as an interesting adjunct to his core belief structure.

For example, my belief structure is rooted in trend following for reasons outlined throughout this book. I believe it is the only methodology that will guarantee the investor or trader to be on the correct side of a big movement in the price, and it is this aspect to trend following that is the key to it being my preference. You see, I am self-aware enough to realise that the thing I hate most in the markets is *not* being involved in a market that makes a big movement in price! How many times have we watched copper, gold, oil, soybeans, wheat, Google, the Chinese stock market and all the other raging bull (or bear!) markets that occur frequently and thought, "if only I'd had a position in those markets"? We all do it. Well, by using a trend-following methodology I can take that thought out of the equation and I can sleep at night knowing that whatever the markets do I will be involved. For me this is important, but for other people it might not be as important and so they will have a different investment or trading philosophy.

However, while my core belief structure revolves around trend following I incorporate other techniques such as volatility analysis, like trading regime analysis, into my process. Essentially, however, I am a trend follower and this process will dominate my investment decisions and the markets in which to trade. In a utopian world I would ideally like to have a trend-following risk allocation to every freely traded market and share in the world, but that is virtually impossible unless you have quite

a few billion dollars in the bank to cover the margin! In the real world, therefore, I like to allocate my trend-following risk to markets where big trends have occurred in the past and, I therefore assume, might happen in the future. Not only this, but I will allocate larger risk to markets that are displaying the characteristics of changing from a range-trading regime into a trending regime. In this way I am using trading regime analysis to choose the markets in which I am to trade, or at least to allocate more capital.

What if someone is not a trend follower? What if someone is self-aware and knows that rather than follow the trend, he likes to take advantage of the trend exhaustion or mean reversion in price? Well, trading regime analysis allows people to choose the markets that best suit their style of trading. There are some markets that can go through years of general range trading and, for our mean reverter, these markets are the ones that he will want to concentrate his efforts on. By utilising trading regime analysis he can focus in on the periods when these markets will be entering a trend exhaustion phase and, hopefully, take advantage of that using certain techniques.

POSITION SIZING

Position sizing is quite simply the most important part of the trading or investment process because it is the variable that has the most potential to make or break performance. If I had bought 100 Google shares, then I am going to have a very different economic outcome than had I bought 1,000 Google shares. It's just so obvious, but it is probably because it is so obvious that the subject of money management or position sizing gets ignored. There are hardly any books on the subject of position sizing and the only two that spring to my mind are *The Trading Game* by Ryan Jones and *The New Money Management* by Ralph Vince. Both are well worth reading in my opinion. Apart from these books there is a dearth of literature on the subject and, even in college or university level investment courses, there is hardly any mention of the "eight hundred pound gorilla in the room". As Gibbons Burke is quoted as saying, "Money management is like sex: Everyone does it, one way or another, but not many like to talk about it and some do it better than others. But there's a big difference: Sex sites on the web proliferate, while sites devoted to the art and science of money management are somewhat difficult to find."

Great traders or investors have the ability to adjust their position size to suit the situation and, in fact, it is probably true to say that some of the greatest traders or investors assumed those mantles because of only a few trades or investments. The point is that these people have made many thousands of trading or investment decisions but, on a few occasions – when they were right about the market direction – rather than make an average size bet, they bet the ranch. I was watching the remake of *Oceans 11* when, in one scene, Danny Ocean played by George Clooney came up with a perfect way to describe this ability and its effect when, as he was giving reasons to Brad Pitt's character for doing the heist on not one, but three, casinos,

he says “Because the house always wins. Play long enough, the house takes you. Unless, when that perfect hand comes along, you bet big. Then *you* take the house.”

Utter genius, and this description of a betting strategy explains why such people as George Soros are such masters. They get it wrong as much as the rest of us but when they get it right they are not afraid to really push the boat out and put on a massive position. When the UK pulled out of the European Exchange Rate Mechanism in 1992 after the Pound fell dramatically against the Deutschmark and other European currencies, George Soros was blamed for being the man who “broke the Bank of England”. All Soros did, however, was see a situation with a highly probable outcome and was not afraid therefore to put on a huge short position in the UK Pound. That one trade made him much more famous than he ever had been previously, but not much is made of the fact that the returns from that trade would have been disproportionately large compared with all his other trades. When Stanley Druckenmiller worked with Soros on his Quantum Fund he tells of the time when he had taken a large short position in the US dollar against the German mark. “How big a position do you have?” Soros asked him. “One billion dollars” Druckenmiller replied (one billion dollars was a massive position in relation to the size of the fund). “You call that a position?” Soros exclaimed (*Soros, The Unauthorised Biography* by Robert Slater). Soros persuaded him to double his position and Druckenmiller later said that the thing he learned most from Soros was the ability to feel comfortable taking the big risk or “to go for the jugular” when you felt it was right. All the great traders have had this ability to feel comfortable with risk and to increase their position size to a level where it makes a real difference to performance.

Goldman Sachs are well known for having a large risk-taking culture and a story from there helps to explain why their proprietary trading teams regularly (but by no means always) perform extremely well. Christian Siva-Jothy tells of the time when he first joined the proprietary trading desk at Goldmans having come from another bank and how he put on his first positions there. He went short 50 million Deutschmarks against the Swiss franc, which was the biggest position he had ever put on in his career. The head of trading came up to him and asked him if he had any positions on and, when Siva-Jothy proudly told him he was “short 50 Mark Swiss”, the head of trading was purported to have said “I like a guy who averages into his positions” (*Inside the House of Money* by Steven Drobny).

Of course not every time we increase our position size will it turn out to be a good thing, but unless we vary our position size we are effectively assuming that every trade has the same probability of success. Bertrand Russell, the British author, mathematician and philosopher, once said that “when one admits that nothing is certain one must, I think, also admit that some things are much more nearly certain than others”, and I believe this philosophy applies to trading. Of course nothing is certain in the markets, of course no one can possibly know what will happen in the future, but there are clues and hints that add to or take away from the probability of a market doing something that means that it no longer becomes like the 50% probability of a coin toss. And that is just about the probability of market direction.

In trading regime analysis we are not, in the first instance, worried about the direction of the market as our primary goal is only to identify whether the market is more likely to trade in a range or have a strong trend. If, for example, we are using a trend-following type of strategy and our analysis tells us that the market in question could be about to enter a period of higher volatility, as defined by a higher standard deviation of the price series (i.e. a trend), then it would make sense for us to add a little more risk to the directional positions we might put on in that market. Conversely, if our analysis tells us that the market in question has already had a big trend and the probabilities favour a period of range trading, then it makes sense to lower our position size.

As an example of how varying the position size depending on a study of volatility can help us, let us consider a very simple example using a long-term trend-following method where we use weekly time fractal moving averages to determine the direction of our position and we refer to the monthly time fractal for an analysis of volatility and, therefore, the likely trading regime. We can assume for simplicity that we will only refer to one indicator of trading regime, whereas in reality we would look for confirmation from other indicators. Figure 23.1 shows the bar chart of the EURUSD exchange rate using the monthly time fractal and a 20-month Bollinger Band Width in the lower chart. This is our reference for potential trading regime because we are actually investing using one time fractal down and we can see from the chart that there were periods of low volatility indicating range trading, and periods of higher volatility indicating trends.




Figure 23.1 EURUSD monthly and Bollinger Band Width

Date	Action	Volatility	Size	Price	Variable P&L	Constant P&L
27/12/1996	Short	Low	€4m	1.2415		
07/11/1997	Long	High	€1m	1.1556	297,335	148,667
26/12/1997	Short	High	€1m	1.1149	-36,506	-73,011
15/05/1998	Long	High	€1m	1.103	10,789	21,578
26/06/1998	Short	High	€1m	1.0942	-8,042	-16,085
07/08/1998	Long	Low	€4m	1.1088	-13,167	-26,335
12/02/1999	Short	Low	€4m	1.1308	77,821	38,911
29/12/2000	Long	High	€1m	0.9427	798,133	399,067
13/04/2001	Short	High	€1m	0.8882	-61,360	-122,720
24/08/2001	Long	High/ Falling	€1m	0.9142	-28,440	-56,880
25/01/2002	Short	Low	€4m	0.865	-56,879	-113,757
03/05/2002	Long	Low	€4m	0.9173	-228,061	-114,030
07/05/2004	Short	High	€1m	1.1883	912,228	456,114
					1,663,850	541,517

Figure 23.3 EURUSD trend following – variable versus constant position size

position size is kept at a constant €2m returned €541,517 over the period, whereas when the position size is varied in accordance with the volatility or trading regime analysis, the return increases to €1,663,850, over three times higher! This shows the incredible power of a volatility-adjusted position size for a trend-following methodology and it shows that the basic concept can reflect reality. When volatility is high, the chances of being whipsawed around are much higher than when volatility is low. For trend followers the goal is to maximise the minority winning trades and minimise the majority losing trades, and this data certainly points to the fact that this is achievable by varying position size based on trading regime analysis.

CHOICE OF TRADING TECHNIQUE

Although everyone has an inherent bias within them of either being a seller of volatility (and therefore being more comfortable with a mean-reverting style of trading) or a buyer of volatility (and therefore being more comfortable with a trend-following or volatility breakout style of trading), there are some traders and investors who can mix and match their trading techniques depending on the situation at hand. This is where trading regime analysis can lend a helping hand because it skews the odds in favour of one type of environment or the other.




Figure 23.2 EURUSD weekly

If we now look at Figure 23.2, which shows the weekly time fractal of the EURUSD exchange rate, we can see our trend-following trading signals. These are given by the interaction of the 5- and 40-week moving averages and we assume a very simple system whereby we cover shorts/go long when the 5-week average crosses the 40 from below and we sell longs/go short when the 5-week crosses the 40 week from above. As such the system is always in the market and is, therefore, continuous.

The position sizing procedures for the system are also very simple and we assume that, when a trading signal is given and the Bollinger Band Width is high, the position size is €1 million nominal, but when a trading signal is given and the Bollinger Band Width is low the position size increases to €4 million nominal. The trading signals and results are shown in Figure 23.3, which shows the date of the trading signal to cover shorts/go long or sell longs/go short, the volatility analysis as indicated by the monthly band width, the position size, the spot price at entry/exit and the profit and loss. The profit and loss calculations are based on the spot return only so no account has been taken of the interest rate differential or carry between the two currencies. Notice, too, that there are two columns for the P&L. The Variable P&L column reflects the profit or loss based on the position size varying depending on the level of volatility, whereas the Constant P&L column reflects the profit or loss based on the position size remaining at a constant €2m.

The data is therefore comparing a situation where the position size of a trend-following system is halved when volatility is considered to be high, but doubled when position size is considered to be low. The figures show that the system where the

Figure 23.4 shows the bar chart of Conoco Phillips, the integrated energy company, using the weekly time fractal from 2003 until 2006 with the Bollinger Band Width in the lower chart. We can see that the market had been in a solid uptrend, pausing for a while as the band width became quite high in both November 2004 and March 2005, and in September 2005 the band width had once again reached quite high levels. At this point a trader in Conoco Phillips might well think that, as the volatility was quite high as measured by the high band width, and that therefore the probabilities of a regime switch away from a trending environment to perhaps one where the market experiences range trading more, the time would be right to trade with a mean-reverting bias. There are many ways to achieve this of course and a lot of people will use discretionary methods of range trading by identifying actual previous price highs and lows, buying at the lows and selling at the highs. This sounds like a great way to make money doesn't it? Buy low and sell high. Unfortunately it is yet another market phrase, which is brandished about by people, that actually goes a long way to cement the bias we have as humans to think that something that has fallen in price must be considered cheap and something that has risen in price must be considered expensive. Buying low and selling high is a good strategy but only if the market in question is in a range-trading or mean-reverting environment.




Figure 23.4 Conoco Phillips weekly and Bollinger Band Width

to take advantage of this regime. Using the slow stochastic as a guide for buying when the market was oversold, and selling when the market was overbought, would have meant that a trader or investor following this process would have outperformed a buy and hold investor or an investor using a trend-following process.

RISK BUDGETING AND STYLE DIVERSIFICATION

Although hedge funds have been around since the 1940s there has been a tremendous growth in their popularity since the mid-1990s and, at time of writing, the absolute return philosophy that is the bedrock of hedge funds is now well entrenched into the mainstream investment community. Changes have therefore been taking place in mainstream investment philosophy to reflect some of the nuances that had traditionally been confined to the hedge fund community. One of these changes concerns the allocation of risk to different methods or different styles of investment.

Style diversification, although it is thought of as a relatively new concept, has actually been part of mainstream investment philosophy for quite a number of decades, even though it may not have been formalised as such. In the stock markets there is a general split between what is known as a value style of investment and what is known as a growth style of investment. The value style of investment means that the focus for a fund's money is to invest in companies where the share price appears to be undervalued by some form of fundamental analysis (for example, price to book ratio or dividend yield). On the other hand, the growth style of investment means that the focus for a fund's money is to invest in companies that exhibit signs of above-average growth in their profits or markets, even if the share price appears to be overvalued by the forms of fundamental analysis like those described above. Investment firms generally used to be grouped into one of these two styles and, while some still follow just one form of investment methodology, most investment firms will have products that will provide exposure to each style.

More rarely will we find single funds that split their styles between growth and value but they are becoming more common, and once a fund has two or more different styles to choose from, the allocation of capital to those styles becomes a crucial variable to the performance outcome. Sometimes an investment company will offer what is known as "fund of funds" products, which are single fund products, but rather than invest directly in the shares of companies it invests in other funds, and one of the aims of doing this is to diversify the exposure to different styles of investment. These fund of funds face the same question. How do they allocate capital between the different investment styles?

One way of allocating capital to investment styles is to use an approach based on portfolio theory and some sort of optimisation technique. This is the way most funds would benchmark their asset allocation to, for example, different geographical markets or between stocks, bonds and cash. An optimal allocation based on historical performance would be calculated and the fund would either retain this weighting

Otherwise it pays to actually employ a strategy of buy high, sell higher or if you are shorting, sell low, buy lower! However, buying something that has already gone up in price is very hard for the human mind to feel comfortable with.

Getting back to our trader in Conoco Phillips, he has decided that since the odds favour a range-trading regime he will base his trading decisions using a mean-reverting technique and the analysis technique he likes best is a slow stochastic. Now since the volatility analysis is based on the weekly time fractal our trader would be best to use one time fractal down when actually trading in the market as the lower time fractal will enable a more timely response to the changes in volatility.

Figure 23.5 shows the daily chart of Conoco Phillips with a slow stochastic from August 2005 until February 2007. Slow stochastic is a traditional method of measuring whether a market is overbought or oversold and, like all oscillators, it works extremely well in range-trading regimes and really badly in trending regimes. It is based on the observation that as prices increase, closing prices tend to be towards the upper end of the price range, and vice versa when prices are falling. A stochastic process will start to level off if a time series shows an uptrend where prices are closing towards the highs but then the closing prices start to close nearer the middle or bottom of the range on average. In this way it measures the underlying strength of the move.




Figure 23.5 Conoco Phillips daily and slow stochastics

We can see from the chart that the high Bollinger Band Width on the weekly time fractal did, indeed, lead to a sideways trading regime for quite a number of months and that a daily time fractal stochastic process would have been well suited

on a static basis or look to adjust the exposure on an “underweight, overweight” basis but within certain allowable parameters. If the fund is retaining exposure on a static basis then the question of how to change allocation does not come up, but if the fund is dynamically changing the capital allocation then the process of deciding how or why to change the weighting assumes an important role. Most funds do this by employing a discretionary process of asset allocation between asset classes and geographical markets based on an analysis of the global macro-economic situation and related topics.

However, if the decision of capital or asset allocation is not one of asset class or geographical markets but one of investment style, then the philosophy is essentially the same, but the analysis would have to change because an analysis of global macro-economic analysis might not be relevant to such a decision. What do become relevant are the *volatility* conditions under which markets are trading because it is these volatility conditions that determine the success or failure of an investment style. This is where trading regime analysis can help.

It is my contention that, whatever strategy you employ in trying to make money from investing in or trading the markets, the success or failure of the strategy will be directly affected by the volatility conditions of the markets. Therefore, any investment or trading strategy can be thought of as either a “long volatility” strategy (that benefits from trending behaviour in the markets) or a “short volatility” strategy (that benefits from general range-trading behaviour in the markets). There is, in the end, nothing else. No strategy makes money consistently in both periods. Ahhh . . . I hear you say. A discretionary investment manager or trader will shift the investment style depending on whether they think the markets will be volatile or not volatile. Well this is the point of the book. Most people do this, if they do it at all, by a reactionary process of guesswork and “feel”. What trading regime analysis can add is a more formalised process of deciding the different volatility conditions and allowing the trader or investment manager to be proactive in shifting investment styles.

Many funds are now allocating capital between different styles that can very generally be split into those that benefit from trending markets and those that benefit from range-trading markets. Investment styles can be split like this because in general they will do well under certain conditions of volatility (as defined by range expansion or range contraction). For instance, as we have just discussed, value investing where there is a return based on, for example, an attractive dividend yield will probably outperform growth investing when the stock market in general is range trading and vice versa when the stock market is trending up. What about other investment styles? Different asset classes like equities, bonds, currencies and commodities will have some different individual styles of investment but, in general, the returns from these investment styles will be to a greater or lesser extent affected by the volatility conditions in the market. Let us now consider a few investment styles and the volatility conditions or trading regimes from which they can benefit.

CARRY STRATEGIES

A very popular investment style, especially in the currency markets, is what is generally known as “carry”. When a currency fund manager says they have a capital allocation to “carry”, what they mean is that they have an allocation to an investment style where the returns are mostly generated by the yield or interest rate differential between the currency or currencies that are sold and those that are bought. This is the famous “currency carry trade” that has hit the headlines in recent years as it has become ever more popular. It involves shorting a currency (or a basket of currencies) that has a low yield and buying a currency (or a basket of currencies) that has a high yield.

Investing in anything has an associated cost. If I invest my cash in the stock market I am actually doing two things or making two decisions. I am investing in the stock market but at the same time I am also deciding not to invest in the cash market, and so the act of investing in the stock market involves a relative cost of giving up the returns I would have earned in the cash market. So, being a relative cost, it will sometimes be negative (when the dividend yield in the stock market is lower than the cash market) and sometimes be positive (when the dividend yield in the stock market is higher than the cash market). Thus, in any investment there is a cost to “carry” that investment because you are making a *relative* decision. There is always an associated cost either in the form of an opportunity cost (as above) or in a real form by borrowing something to invest in something else. This cost is the cost of carry.

In effect, the currency fund manager, by shorting the low yielding currencies, is borrowing them (he is shorting something he does not own) and so the cost of borrowing these currencies is the relevant domestic interest rate. However, he is using these borrowings in low-yielding currencies to invest in the high-yielding currencies on which he earns the domestic interest rate on those currencies. Therefore he earns the interest rate differential (the carry) between the low yielders and the high yielders. The currency fund manager is seeking to benefit from the positive cost of carry. If he had shorted the high yielders and invested in the low yielders, he would have a negative cost of carry because he would be borrowing something with high interest rates to invest in something with low interest rates.

This, on the face of it, sounds extremely attractive. We can borrow something with low interest rates and invest in something with higher interest rates and rake in the differential. It's like free money! Oh, if it were that simple. Unfortunately there is another part to the story. In both the stock market versus cash returns example and the currency carry trade strategy there is another variable element to the returns that these strategies will generate and that is, of course, the capital performance of the stock market and the performance of the spot exchange rate. These elements will be affected by the *volatility* conditions in the market and there is the obvious conclusion that, when volatility is low and we have a range-trading regime, the positive carry earned will be attractive whereas, when volatility is high and we have a market where

the spot returns are unstable, the positive carry earned might be completely nullified by a negative return on the capital side of the equation.

Carry strategies also exist in single country markets. In fact this is where the strategy is most heavily used, and this is by bond or interest rate market participants very simply borrowing at low interest rates and lending at higher interest rates. This is how banks make their money and it has been the staple strategy in interest rate markets for centuries. However, this is not an arbitrage and there are always risks. If someone was borrowing short-term money at low interest rates and lending long-term money at higher interest rates, then they face the risk of high volatility in the yield curve and the bond market. What they want is for the yield curve and bond markets to be stable to enable them to earn the “carry” on borrowing short and lending long but, if they are not stable, the carry earned could be wiped out, and more, by losses on the capital side of the equation.

Therefore, the carry or interest rate differential strategy goes through periods of working extremely well and other periods of working extremely badly. The periods when it works well generally tend to be low-volatility regimes and the periods when they work badly tend to be periods of high-volatility regimes. Of course, a carry strategy can work magnificently when there is an interest differential to be earned *and* the capital side of the equation is in a trending (high-volatility) regime in the direction of the investment (the low-yielding capital side is going down relative to the high-yielding capital side). In these cases, the carry investor gets an extra return, but the risk is still present that the capital side could have gone the other way and wiped out the carry. Therefore, a trend can be the carry investor’s friend, but only in one direction. High volatility, in general, remains the carry investor’s enemy.

OPTION STRATEGIES

Taking Advantage of Trend Exhaustion

Trend exhaustion is a generic term that means those levels or points where the trend is so mature it is coming to an end. In my opinion it is very important to realise that this does not necessarily mean that the trend has to reverse direction and start a new trend in the opposite direction. When I was a novice fund manager looking at oscillating type indicators that indicate levels of a market considered to be overbought and oversold, I made the mistake of assuming that these overbought or oversold conditions would mean that the market had to go in the opposite direction before the indicator could return to more neutral levels. Perhaps I should have looked at the specifics of the calculations of these indicators a little more because it is a mathematical certainty that, for most oscillator indicators, levels of overbought or oversold can be “worked off” by the market merely going sideways.

What this means in practice is that when something is overbought or oversold, or better still, showing trend exhaustion based on some volatility analysis, then it

is perhaps just as likely to trade sideways as it is to reverse direction and trade in the opposite direction of the previous trend. So surely it would make sense for market traders, who are looking to profit from this likelihood, to use a strategy that benefits if the market trades sideways or against the previous trend. If we think about it, a market, over any time period, can have five different outcomes. It can go up a lot, go up a little, remain unchanged, go down a little or go down a lot. Now let us suppose that we are bullish of the market in question and think that it will go up. We think it will go up but, of course, we can never be sure. If we buy the underlying market in the traditional manner via stock or futures, then we benefit if the market goes up a lot or if it only goes up a little.

That means that we benefit from two out of the five potential outcomes or, put another way, we have a 40% chance of success. Now let us suppose that, rather than buy the stock or futures market we sell an out-of-the-money put option. What this means is that we are selling an option below the current market price where we are accepting to take delivery of the market or stock at that price. If the market fell below that strike price then the buyer of the option would be able to sell or “put” the stock or market to us at the strike price, meaning that we would be disadvantaged. Therefore, by selling the put option below the current market price we are implicitly saying we are more bullish than bearish of the market because we are willing to write an option where we are disadvantaged if the market falls a long way. However, the crucial point here is that we are not limiting ourselves to benefiting only if the market goes up. By selling the out-of-the-money put option we benefit if the market goes up a lot, goes up a little or stays the same and, because we have sold an out-of-the-money put option, we even benefit if the market falls only a little. So we benefit from four out the five potential outcomes, which is an 80% chance of success. Naturally, if we have an 80% chance of success in the strategy then we should realise that the potential reward is going to be much lower than a strategy that has a lower chance of success. This simply has to be the case.

When we are anticipating a drop in volatility it doesn't have to mean that the underlying market will go into a very narrow range-trading mode. It often does, but it doesn't have to mean that. Volatility as measured by the standard deviation can drop with the underlying market moving in what looks like a sharp (trend-like) move in the opposite direction. Of course, if this move turned into a sustained new trend then the standard deviation would continue to move up, but at least the initial bit of a trend exhausting itself can seem to be a sharp move. Nevertheless, when the standard deviation has been high relative to its history, then the odds are skewed in favour of some sort of choppy price action developing.

Taking advantage of this price action can be difficult. One way is to sell option strangles (which we will discuss later) but this type of strategy for selling volatility is more suited to a very distinct range-trading regime. Another way of taking advantage of trend exhaustion and the anticipated drop in volatility is to be a little more directional in an options strategy. This can be done in various ways but the most obvious would be to sell a naked out-of-the-money option, either a call in the

case of an uptrend exhausting or a put in the case of a downtrend exhausting. This can be a good strategy but it is *very risky* as you are short an option with no backing (as in the case of a covered call write) and so the downside is unlimited. Not only that, but because the strategy involves selling an out-of-the-money option that will have a low delta, it will however have a high gamma if the market starts to go against you. That is, the rate of change of the delta will be very, very fast if the position starts to go against you and this, as anyone who has been short naked options knows, can be very destructive to both your fund and your emotions!

Let's look at an example using the euro exchange rate with the US dollar. On day one you notice that the weekly standard deviation in the EURUSD rate is at the higher end of its history and looks like it might be turning over. You anticipate that the recent uptrend may be exhausting and you want to take advantage of this drop in volatility. You think the trend may be exhausting given the high level of standard deviation, but you know that trend exhaustion does not necessarily mean a sustained reversal of the previous trend and could just as easily mean a sideways choppy price action for a while before the previous trend reasserts itself. In order to take advantage of this you decide to sell a call option rather than short the cash market, because by selling the option you will benefit if the underlying market goes sideways or down, whereas by selling the cash market you only benefit if the underlying market goes down. The market is trading at 1.4665 and you decide to sell the 1.5000 call for one month's expiry. You write the option in a nominal amount of €2m and, for this, you receive €5,793 in premium. This call option has a delta of 20, meaning that there is a 20% chance that the option will end up in the money at expiry. In theory this means that there is an 80% chance that the option will not end up in the money and, because you have sold the option, there is an 80% chance that the strategy will be successful. This sounds wonderful. You have just deposited almost €6,000 in your account, and there is an 80% chance that the strategy will be successful and you will be able to keep the cash. However, the 20% chance that the strategy is not successful is not as sanguine as it sounds, because when this 20% probability turns into reality things can start to go wrong very quickly.

Let us assume that the market does not, in fact, start to trade sideways or down but continues its upward trend. Let us assume that two weeks after the trade was put on, the market has moved up and is trading at 1.5000. The call option you sold still has two weeks to run to expiry and yet the underlying market is trading at the strike price already. The delta on the option (the probability that it will expire in the money) is now 50 and the value of the option is €18,247. If you wanted to cover the position now and buy the option back, it would cost that amount – meaning that you would have to pay out €18,000! You have received €6,000 and so this results in a net loss of €12,000 or 200% of the maximum possible gain from the strategy. If you held on, thinking the market would come back but instead the market continued to rally, then it gets even worse at an exponentially higher rate. If, one week before expiry, the underlying market is trading at 1.53 then the call option with a 1.50 strike now

has a delta of 76 and would cost around €47,000 to cover, a net loss of around 683% of potential maximum gain!

This sounds disastrous (and it is) but, surprisingly, it is not as bad as if you had decided to sell the underlying market rather than write the call option. If you sold the underlying market at the cash exchange rate of the time then, under the scenario where the market is trading at 1.53 with one week to expiry, you would have lost €73,000 on your €2m nominal position. Now, we could say that, because the option we sold had a delta of 20 at the time of initiation, then we should really compare the strategies with an equivalent cash position. In this case it would be €2m multiplied by 0.2, which is €400,000. That is, the amount of an equivalent cash position would have been €400,000 nominal and if this was the case then the cash market loss would be €17,000, a lot less than the short option strategy. There is some sense in comparing the two strategies based on this delta-adjusted nominal cash position, and if we do this then the options strategy looks much more risky. Had you sold the cash market and the rate had fallen, then you would gain far, far more than the short call option trade but this is the trade-off you make. Selling naked options is considered to be very, very risky and most people would rather chew their own foot off than go short naked options. This is because the delta effect, as seen in the example, works against a short option position very, very quickly. Being short gamma can be an extremely stressful position to be in.

Selling options, either naked or via a strategy that has a net negative gamma effect, is not for everybody. I think you have to have a certain personality to be comfortable living with a short gamma strategy because, in the case of trend exhaustion, it essentially requires you to “stand in front of the train”. Now, the point of trading regime analysis with regard to identifying trend exhaustion is to identify when we think that train is slowing down and possibly reversing. Nevertheless, to take advantage of the trend exhaustion still requires us to stand in front of the train, and even a slow-moving train can do some damage if we have a negative gamma strategy. When selling options at trend exhaustion points the trader has to be comfortable with a risk-to-reward profile, that means he will risk three or four times the potential reward. For some people, and I include myself, that risk-to-reward profile just does not fit with their personality. Some people are natural buyers of volatility and some people are natural sellers of volatility. Nevertheless, selling volatility at trend exhaustion points can be a profitable strategy if you can manage the short gamma risk appropriately.

Beefy versus Boycott

It is your personality that will go a long way to deciding whether you are more comfortable employing a long-volatility strategy or a short volatility strategy, and another sports analogy can be drawn here. Cricket is a fine game and Ian Botham (known as Beefy due to his well-built frame), the English all rounder, was a magnificent exponent of his craft. He was a great bowler and when he was batting he would tend

to play with panache by attempting risky shots to the boundaries. Another English cricketer, Geoffrey Boycott, however employed a much more diligent approach to scoring runs by “nudging singles” to build a score and his innings were well known for being long, and some would say boring. Now cricket aficionados could point to the fact that one of the reasons for this difference in batting style was because of the fact that Boycott tended to open the batting and his job was to build a slow steady score, whereas Botham was a middle order batsman and he was expected to go for quicker runs. There is undoubtedly some truth in this but the fact remains that both players’ personalities suited their style of batting. In market terms Botham was a natural buyer of volatility whereby the risk (of getting out) was much lower than the reward (hitting a six). Boycott, on the other hand, was a natural seller of volatility whereby the risk (of getting out) was arguably greater than the reward (hitting singles). If we were talking about baseball rather than cricket, we could replace nudging singles for bunting, and hitting sixes for “swinging for the fences”. Babe Ruth, who held the record for the most home runs, was a prime example of a natural buyer of volatility and this was encapsulated in his famous quote that “every strike (out) brings me closer to the next home run”. The Bambino was not a natural bunter!

Those who are not too comfortable with an inverse risk to reward equation that comes from selling options at trend exhaustion points, could employ a strategy that has a positive risk to reward but can also (sometimes) benefit from a sideways range-trading movement in price. If a market has fallen to the extent that the volatility indicators and other trading regime techniques suggest that trend exhaustion is setting in and that a reversal and/or period of range trading is set to take place, then, rather than write out-of-the-money put options or establish a bull put spread for a net credit, the investor or trader could employ a bull call spread.

The bull call spread involves buying a close-to-the-money call option and, at the same time, selling an out-of-the-money call option. The rationale behind the option structure is that you expect the market to go up, but probably by only a little. In this way the structure can benefit from a sideways to up movement in the price that would normally come from a downtrend exhaustion point. The bull call spread is similar to the bull put spread, but there are subtle differences. The bull put spread, being net short theta, has a slightly higher probability of success than the bull call spread. However, to take account of this the bull put spread has a less attractive risk-to-reward ratio than the bull call spread. The same analysis applies for bear put spreads versus bear call spreads.

Option Strangle Selling

One of the most popular ways to take advantage of a range-trading regime in a market is to trade what is known as a strangle strategy. The options market is a fantastic market for structuring trades around a view that you may have regarding

the underlying price movement, and in the case of strangles (as well as its close cousin the straddle) the view will be regarding the volatility of price action.

There are two types of strangle strategies, long and short. A long strangle strategy involves the purchase of an out-of-the-money call and a purchase of an out-of-the-money put and so the strategy benefits if there is a big price movement either way. The person putting this strategy on is expecting a big price movement but is unsure of the direction in which the price movement may resolve.

A short strangle strategy, on the other hand, involves a view that the market will stay in a range and involves the sale of an out-of-the-money call and a sale of an out-of-the-money put. This strategy will benefit if volatility remains low and the market stays in a range.

We can see clearly from the pay-off diagrams in Figures 23.6 and 23.7 how each strategy benefits from the aforementioned movement or non-movement in price.




Figure 23.6 Long strangle at expiry pay-off

Figure 23.6 shows the pay-off at the expiry of the options for the long strangle strategy. The horizontal axis shows the price of the underlying instrument and the vertical axis shows the profit or loss on the options position. We can see that in this strategy the maximum loss is limited at point L, this being the amount of premium paid for the options, and we can see that the maximum profit is unlimited because, being long of the options, we can benefit from extreme price movements. The price of the underlying market is at 50 when we put the trade on, and the strategy makes money if, at expiry of the options, the price of the underlying is above 55 or below 45.

Figure 23.7, on the other hand, shows the pay-off at the expiry of the options for the short strangle strategy. We can see here that it is the profit that is limited to point P, this being the amount of premium that is credited to the account by the sale of the options, and we can also see that the maximum loss is unlimited because, being short of the options, we are at risk from extreme price movements. The price of the underlying market is at 50 when we put the trade on, and the strategy makes money if, at expiry of the options, the price of the underlying is below 55 and above 45.




Figure 23.7 Short strangle at expiry pay-off

It is clear, then, that the long strangle strategy benefits from extreme price movements in the underlying market and the short strangle strategy benefits from an underlying market that trades in a range. In essence, therefore, the strangle strategy is a play on underlying market volatility, which makes it an ideal candidate to trade an analysis of the trading regime of the underlying market.

In fact many fund managers have a capital allocation to a short strangle strategy and are, in effect, passively selling volatility all the time. Many people believe that selling volatility is a good long-term strategy because the markets will generally have more periods of low volatility than high volatility. This is probably true if we think of the fact that a trend-following strategy will have a profitable trade ratio of less than 50% (less than 50% of trades are winners). Therefore ranging behaviour tends to dominate the markets, but please note that this does not mean that trend following does not make money over the long term. However, due to this fact, it is certainly true to say that selling options can, on balance, be a solid strategy to use if you are aware of the risks involved.

An option can be thought of as an insurance policy. One of the most ironic aspects to the markets is that great deals of participants are still wary when the word option is mentioned. People on the sell side (the brokers or bankers who are acting as intermediaries) who try to sell option ideas and trades to the buy side (fund managers or corporate treasurers who have risk to manage) are continually frustrated by the fact that many buy side people still regard options as dangerous and volatile instruments. An ex-colleague who used to be on the sell side of the industry tells some good stories of how he would try to persuade corporate treasurers to use currency options to mitigate their risk. He would ask the person in question if he had a house, a car or life insurance policies and the person would generally answer "Of course I do, it would be stupid not to have house, car or life insurance in case an accident or anything bad happens." Then the trap was set. "Why then", my friend would ask, "would you never consider using an option to mitigate your foreign exchange risk because having an insurance policy on your house, car or life is in effect an option contract

where you pay a premium for protection in case, as you say, anything bad happens!” Usually there would be silence on the end of the phone as the penny dropped but the majority of time the corporate treasurer would still be wary of buying the option for protection. Deciding not to *buy* the option, however, would arguably be the correct course of action to take.

Options are like insurance policies, but let us ask ourselves two questions: Who benefits over the long term from the insurance industry? Is it the customers like ourselves who buy the insurance policies to protect ourselves from accidents or is it the insurance companies that sell these policies to us? Of course it is the insurance companies because we all know that the amount of times we claim on our insurance policies is tiny compared to how long we have been insuring ourselves and paying the monthly or annual premium to do so. Insurance companies make their money by selling us options that we can exercise if something bad happens. Essentially what they are doing is selling volatility because they are assuming that nothing bad will happen the majority of the time. If things start to go bad a lot more often (for example, if there is a spate of natural disasters), then the premium for insurance will go up and this is just the same as if volatility in the markets rises, then the premium for options will go up.

It is therefore the sellers of options who will generally make money most of the time but, just like insurance companies, when the big accident comes along, there will sometimes be a big loss to take. One of the main reasons that option sellers make money most of the time is due to the “theta” effect, or the time value effect. This relates to the fact that as an option nears the expiry date of the contract, then the time value has an ever-increasing detrimental effect on the overall value of that option because, simply, the time left for “something bad to happen” is running out quickly.

Figure 23.8 shows the extremely compelling case for being a seller of options rather than a buyer. It shows that as the expiration date of the option gets closer, the time value of the option goes down but, crucially, it does so at an ever-increasing pace. In this regard the seller of the option is gaining theta (time value) every day but the closer the option comes to expiry that daily gain becomes higher and higher. As each day passes with the price not changing much, the seller of the out-of-the-money option is gaining at an increasing rate. On the other side, the buyer of the option (who wants volatility/movement) has to rely on a bigger and bigger movement in the underlying in order to gain any sort of value from the option because the time value is dwindling at this faster rate.

In this regard sellers of options also act like bookies at the horse racetrack or the house in the casino. Bookies effectively sell options to punters who are going long the option of backing the particular horse at the dealt at odds, and the expiry date is obviously the end of the race. There is a finite time for something to happen, of course, and when the race is in motion the time value of the “options” dives rapidly. Most punters, by the way, act the way most investors do in the options market because at the racetrack we are enticed by the long shots at 30 to 1 odds and in the




Figure 23.8 The theta effect (on an option price)

options market we are enticed by buying out-of-the-money calls or puts with similar big potential pay-offs. This is just another function of human behaviour in that we are naturally drawn to look for the big one-off money-making scheme rather than something that provides a low but steady return. It would seem that what we should be doing in trading options is acting like an insurance company or a bookie. And has anyone ever heard of a poor bookie?

Sellers of options will make a small amount of money the majority of the time but will lose a large amount of money the minority of the time. The analogy that people like to use with regard to option sellers, and indeed any strategy that sells volatility, is that it is “like picking up pennies in front of a steamroller”. Most of the time you will pick up the pennies with no problem at all, but take your mind off the job at hand and not manage your risk appropriately then SPLAT, you get flattened. In fact, most sellers of volatility will get flattened on a regular basis as the volatility cycle ebbs and flows, and it is very difficult for them to manage the risk effectively enough to avoid the odd big loss. Nonetheless, with the important proviso of having robust money management rules, over time this strategy does seem to work on a passive basis and, if this is the case, then surely we can add some value to it by identifying potential SPLAT points (i.e. when the market has potential for a big increase in volatility)? This, of course, is where trading regime or volatility analysis plays its part.

The Case for Long Strangles

Realising that the sellers of options are the people who make money most of the time might lead us to think that we should always be looking to sell options rather than buy them, and there is a good case for this. However, the SPLAT points of a passive option-selling strategy indicate that there will be times when being long options (for example, in the form of a strangle strategy) would be the best thing to do. An analysis

of the potential trading regime should be able to tilt the odds in relation to selling a strangle or buying a strangle, but let's consider the case of just passively buying strangles and see what effect that would have.

As we can see from Figure 23.6 the long strangle strategy is one where the downside is limited to the option premium bought, and the upside is unlimited at least to the upside because the underlying market can potentially go up to infinity and down to 0. If we just passively bought strangles then our profit and loss profile would look like Figure 23.6 over the long term. That is, we would have limited our losses to the small amount of option premium but would still have periods (albeit the minority of the time) when our profits would be large. In fact our average profit would be higher than our average loss and so, over the long term, if we employed this strategy in long-term option contracts, this strategy should also make money.

A good way to think about the effect of having the average win higher than the average loss is to use the coin toss example referred to earlier and reiterated here. Tossing a coin over the long term should produce 50% heads and 50% tails, so in any one coin toss there is a mathematical 50% chance of being right or wrong by calling heads or tails (yes, there can be streaks of heads and streaks of tails but over the long term the average is 50%). Let's say you play a game whereby for each time you call correctly you win £1 and each time you call incorrectly you lose £1. Would you play this game? Given the fact that you will win 50% of the time, over the long term, your expected return is zero because for every pound you win you will also lose a pound. So playing this game is a waste of time over the long term. Now suppose the game changes and for each win you get £2 while for each loss it remains at £1. Even though your win/loss profile stays the same at 50%, would you play this game? In the jargon of the markets you would bite the other player's arms off to play this game because it is a proverbial goldmine. Over the long term you are still losing as many times as you are winning but for each win you receive double what you pay out in losses. If you played this game 1,000 times, you would be £500 in profit, whereas if you played the previous game 1,000 times, you would have zero profit. From this example it is easy to see that the amount of times you win or lose is not nearly as important as the average win when you win, as against the average loss when you lose.

In the previous section on option strangle selling the average win is lower than the average loss but this is compensated for by the amount of winning times being far greater than the amount of losing times. However, robust money management has to be applied to a passive strangle selling programme to limit the losses if it is to win over the long term. If we sold the 10-delta strangles with no money management, then we could expect to win 90% of the time and lose 10% of the time. Relating this to the coin toss example, but with this win/loss profile, let's say we won £1 every time we won, but on the 10% of the time we lost, we lost £9. If we played this game 1,000 times, we would win £900 and lose £900, making our net profit exactly zero. Zero profit, even though we are *correct* 90% of the time!

When we are considering when to sell options and when to buy options via a strangle strategy, the trading regime analysis will definitely be able to give the trader or investor an edge in determining whether volatility is more likely to go up rather than down, but perhaps the most important consideration is the theta effect or the time value effect. When we sell options we are selling a wasting asset, and when we buy options we are buying a wasting asset. In this regard buying options can be thought of as analogous to buying a new motor car. When we drive the car out of the garage forecourt we know that each day we own and drive that car it will be worth less for resale value even if the overall condition of the car is still as good as new. This is because cars have a finite life (and from personal experience some cars have a more finite life than others!). It is the same for options. An option has a finite life, so with each passing day that nothing happens and the market does not move, the option loses a little bit of time value and the owner of that option loses a little bit of value, whereas the seller of the option gains that little bit of value. When an option nears the end of its life the theta effect rises exponentially, simply because time is running out for something to happen. In this regard, therefore, it makes much more sense to sell options that are near the end of their life and, by the same token, buy options that have a lot of time value left.

Therefore, when the trading regime analyst is thinking of selling strangles, he should be looking at short-dated options, but when the strategy is to buy strangles then he should be looking at far-dated options. Being short volatility in the near term and being long volatility in the long term is a generally successful strategy for option market makers, and it makes sense when we think that, in the short term, options lose their value quickly whereas, in the long term, there is much more time for a big market move to occur.

Being long strangles over the long term thus produces a small amount of big winners and a lot of small losing trades. This return profile (and that of a long straddle strategy) is very similar to trend-following return profiles because the concept is the same. Limit your losses and maximise your profits so that you effectively run your winners and cut your losers.

Options strategies are very useful in taking advantage of volatility trading regimes and, because of their flexibility, they can be used in many different ways to structure a strategy that suits the user.

TREND FOLLOWING

Trend following is a staple style in many investment management programmes, especially on the hedge fund side and the CTA (Commodity Trading Adviser) side, but hasn't yet fully been accepted in traditional investment houses like pension and insurance funds.

As described earlier, the investment style of trend following does exactly "what it says on the tin". The philosophy of trend following begins with the assumption

that, due to the fact that no one can consistently forecast anything about the future, forecasting the future is therefore a futile exercise and it is best to simply follow the trend of the market. The market will go along its paths for a vast multitude of “reasons”, and trying to figure out the reasons why the market path was, is or will be going in that direction is irrelevant when compared with the reality of what *is happening* at any particular moment of time. Trend followers live in the cold, hard reality of what the market price *IS* and not what someone thinks it will or should *be*. The big assumption that trend followers make is that a market will, at some point, exhibit trend-like behaviour, and to be guaranteed to be on that trend that market has to be part of a trend-following system.

Trend followers realise that markets do not trend all the time and, in fact, will probably only exhibit strong trends about 30% to 40% of the time, but they do believe in the mathematical certainty that if your 30% winning trades have an average profit that is way above the average loss on your 70% losing trades, then a trend-following method will, over the long term, make money.

Long Term

Most of the very successful trend-following programmes are long term in nature. People like Bill Dunn and John Henry realise that when it comes to trend following it pays to be as long term as possible and there is one very good reason for this. Due to the fact that even a good trend-following system will only have winning trades the minority of the time, the aim has to be to maximise those winning trades relative to your majority losing trades. Due to trading costs (not so relevant in FX but relevant in stocks, bonds and commodities), every time a trade is done there is a slippage from the potential profit and if you are following the very short term trends then the system could have lots of trades to process. Therefore, the trading costs will substantially dilute the success of the system, but if the system is following the long-term trends, the trading costs will have much less of an impact on the performance. Another reason for following the long-term trends as opposed to the short-term trends is that when a trend occurs, the long-term trends will deliver much bigger profit than a short-term system. The ratio between your average win and average loss has the potential to be a lot bigger when following long-term trends because a winning trend could last for months, if not years, whereas the losing trades will be cut short much sooner.

Diversification

One of the least talked about aspects of trend following is that it is essential to have a broadly diversified portfolio on which to follow the trends. This is just common sense for two reasons. Firstly, having a broadly diversified portfolio will mean that, in the overall portfolio, no one particular market or stock will have a dominant effect

on risk. This is what I call the “sleep-ability test”. If your exposure to one market or stock is so big that you are thinking about it when you are trying to sleep, then it has too large an exposure in your portfolio. Secondly, although an individual market will only trend the minority of the time, there will always be a market somewhere that is exhibiting strong trends and hence providing some good performance for your portfolio. By not putting all our eggs in one trend-following basket we can add value by having multiple baskets.

There are many different ways to follow the trends of the markets but the most important element to look for in a trend-following system is the difference between the average win and the average loss. To reiterate, although a trend-following system will lose more times than it wins, it makes its money over the long term by making much more on its winners than it loses on its losers. So if I am going to allocate some capital to a trend-following investment style, the one element I have to see in a system is that although it is losing more times than it wins, the average win is much higher than the average loss. It may seem strange to look for a system that is losing more times than it is winning but, if I am allocating capital to an investment style, I want to make sure that the style will win in certain periods and lose in certain periods. This is called style diversification. If the trend-following system is not showing these two characteristics, then I am immediately suspicious that it is something other than a vanilla trend-following system and therefore not a style to which I want to allocate capital.

Moving averages are probably the most popular way to follow the trend of the market and they can be used in a variety of ways. A very basic method would be to take the difference between the closing price and a moving average. If the price is above the moving average, the system is long; and if the price is below the moving average, the system is short. Another way is to take two moving averages (a long one and a short one) and go long when the short crosses the long from below and go short when the short crosses the long from above. Some systems use this method but with a number of moving averages to scale in and scale out of positions. Yet another way to use moving averages in a trend-following system is to calculate the slope of the moving average, and if it is pointed up the system is long and if it is pointed down the system is short. The calculation of the slope could take the form of, for example, today's moving average level minus the level five days ago. If it is a positive number, the slope of the moving average is deemed to be up, and if it is negative the slope is deemed to be down.

Generally moving average-based trend-following systems are known as “continuous” systems, whereby the system will always have a position, long or short, in the market. They differ from non-continuous systems that were covered in Chapter 15 where the system can go through periods of having no position in the market and therefore are, in effect, an attempt to identify trending and range trading regimes via one system. Non-continuous systems can be both pure price based (like the Turtle system discussed earlier) and moving average based (like the Moving Average Envelopes discussed in Chapter 9).

Volatility Breakout

One of the most consistently successful trading strategies of all time is directly related to trading regime analysis and is, in fact, intertwined in the subject. The venerable volatility breakout strategy is trading regime analysis personified because it aims to take advantage of a change from a period of low volatility to a period of higher volatility. These changes, as we have seen, can produce fantastic profitable opportunities to benefit from strong, trending moves in the markets because the change in psychology from a period of tight range trading to a breakout of that range can and does lead to extended moves in price that feeds on itself.

When the market is trading in a tight range the psychology of market participants becomes ever more complacent. The trading range is produced by mass, net uncertainty and/or neglect and quite often is accompanied by a drop off in volume as the range continues. Market participants either don't care about the market or, if they do, they become used to the range extremes holding, and these range parameters subsequently become burned into their psyche.

When the market eventually breaks out of the ranges these market participants are caught off guard and rush in either to cover their positions (shorts in case of an upward breakout and longs in case of a downward breakout) or to put fresh positions on in order to benefit from the move. It doesn't matter what the news is or what the background fundamentals are – when the range breaks it triggers a psychological reaction in the human market participants which makes them think about the market and, more importantly, their own position in the market. It is this psychological reaction that leads to the extended move because the inner voice of the market participant is screaming out to “*do something*”.

It could be that I went short of the market as it went up to the top end of the trading range, expecting the range to hold again and fall back but if, instead, the market breaks out the top end of the trading range my psychology changes rapidly. My inner voice is telling me to cover the short as quickly as possible because the market price is breaking into new highs over and above the range that we have comfortably been in for the last period. It might be that I have no position in the market and I am not even interested in watching it because of the fact that whenever I did look at the price it was always between 40 and 45. But hang on, I just looked at it and it was trading at 47. I need to look at that market because it might be on the move again and I do not want to miss a big move! I phone my broker and buy at 47. The point is that when a market starts to move away from an established range of prices it triggers deep psychological and behavioural urges within us as human beings and, because we are human, acting on these urges in such a state of mind becomes easy to do. Moreover, these urges can be so powerful as to overwhelm or change, as part of the cognitive dissonance process, any thoughts one may have as to other aspects of market analysis, such as the so-called fundamentals.

Let's take as an example the NZDJPY exchange rate during April and May 2007 (Figure 4.7, reproduced here). The market stops its uptrend in the middle of April




Figure 4.7 (reproduced) NZDJPY daily – classic rectangle

at around 89 the figure. (In the foreign exchange markets the convention when talking about big round numbers is to use the expression “the figure” so as to avoid any element of doubt as to what is being referred to. For instance, in this example, if I was talking to someone and had just referred to “89”, they might have thought that I was meaning a rate of 88.89 (eighty eight spot eighty nine) – that is, the quoted decimal points of 89 rather than the “big figure” of 89.) So the market stops its uptrend around 89 and establishes a horizontal range between 89 and the high 86s low 87s for a number of weeks. Then on 31 May the market breaks up past 89 and does not come back. Instead, the next day it breaks up some more to well above 90. What do you think the psychology of the market participants is at this point? The market has been in a tight trading range for six weeks and suddenly starts to break up. There was no particular “reason” for the break up in price out of the range. In fact, the economic statistics release calendar for 31 May 2007 point to the fact that New Zealand business confidence actually showed a steep fall when it was released that day and that Japanese construction orders and housing starts statistics both came out stronger than expected that day. If you were a fundamental analyst you would have expected the exchange rate to go down (a stronger yen and weaker kiwi). Nevertheless, the market moved up (a weaker yen and stronger kiwi) because, as has been said, the economic statistics are really not as important as the overall supply and demand mix of the market itself, and the supply and demand mix of the market itself may or may not be affected by the economic statistics or the so-called fundamentals at that point of time.

Therefore, the market moves up and there is no “reason” the rationalisers (the fundamentalists) can point to in order to justify it. People who were short of this

the long volatility strategy camp. On one level you could say that value investing is a mean-reverting (and therefore short volatility) strategy because generally the value analyst will look at historical levels of the indicator and assume that the market will revert to some sort of equilibrium from these extreme levels. However, it is the longer term nature of the analysis, which means that in order for that to happen there has to be a good sized move to take the market back to those levels where it could be considered fair value. Of course, the market does not have to do this move all at once in a trend-like fashion, but the nature of the analysis suggests that if the market is extremely over- or undervalued, then the analyst is expecting a trending move back to fair value. This is implicit in the analysis.

Let's take the example of the currency markets valuation techniques. Putting a value on currency relationships is one of the most difficult and frankly futile tasks one can hope to undertake. There are many different methods of valuing a currency, and they range from the very mathematical to the very simple. They can also come up with very different valuation levels and even the same methodology can come up with different values depending on the parameters being used. What, then, I hear you say, is the point of valuing currencies?

Well, for better or worse a measure of value is one of the fundamental aspects that a human being is conditioned to automatically think about when considering an investment, so valuation measures remain a part of currency analysis. There is also a political aspect to it because governments want to know if their currency is under- or overvalued to enable them to reflect this in their economic modelling process and policy response.

As mentioned, valuation models are numerous but the main categories would be under any of Fundamental Equilibrium Exchange Rates (FEER), Behavioural Equilibrium Exchange Rates (BEER), Purchasing Power Parity (PPP) and Real Effective Exchange Rates (REER).

The most well-known valuation technique for currency markets is Purchasing Power Parity, or PPP as it is usually known. The philosophy of PPP is that a good should cost the same in each country or currency block after inflation rates have been taken into account. However, there is a great deal of ambiguity with this such as what measure of inflation to use (consumer prices, producer prices?) and what start date to use in the time series, so there can be various valuation numbers for PPP. Nonetheless it remains the valuation measure that is most widely used in the currency markets and the one that many fundamental analysts and traders base their decisions on. Some people have attempted to make the idea of PPP more readily understandable by using everyday modern products as the reference point. Thus we have the "Starbucks Index" where the price of a tall cappuccino coffee is compared in various countries and the *Economist* magazine has for many years now published the Big Mac Index where the price of that wholesome sandwich offered from McDonald's is compared around the globe. The Big Mac Index, some believe, is a very good reference for a PPP methodology because the ingredients that go into the sandwich will mostly be sourced locally and so should be a true reflection of domestic inflation

rates. Another important tenet of PPP is that the goods being compared should be the same in each country, and so the Big Mac passes this test because it is as homogeneous a product as one could hope to find in various countries around the world. “It just tastes the same as back home” is a well-known expression! The *Economist* Big Mac index is an annual calculation and Figure 23.9 shows the 2007 table.

Country	2007 USD Price of Big Mac Sandwich	Over/UnderValued Versus USD
United States	\$3.41	0.00%
Canada	\$3.68	7.90%
Brazil	\$3.61	5.90%
Mexico	\$2.69	-21.10%
Argentina	\$2.67	-21.70%
Switzerland	\$5.20	52.50%
Denmark	\$5.08	49.00%
Sweden	\$4.86	42.50%
Eurozone	\$4.17	22.30%
United Kingdom	\$4.01	17.60%
Russia	\$2.03	-40.50%
South Korea	\$3.14	-7.90%
Australia	\$2.95	-13.50%
Japan	\$2.29	-32.80%
Thailand	\$1.80	-47.20%
China	\$1.45	-57.50%

Source: *The Economist*

Figure 23.9 Big Mac Index

Figure 23.9 shows the USD-adjusted price for the Big Mac in selected countries and the subsequent percentage over- or undervaluation. We can see that the USD-adjusted price for a Big Mac in Switzerland is \$5.20, and when we compare that with the price in the United States then the conclusion is that the Swiss Big Mac (and therefore the Swiss currency) is 52.50% overvalued against the US dollar. Likewise, if we look at the USD-adjusted price of a Big Mac in Japan we can see that the Japanese yen seems to be undervalued by some 32.80%.

Therefore, based on this information, a value-based currency investor might conclude that the Swiss franc is very overvalued and that the Japanese yen is very undervalued. He therefore might conclude that there is a very good chance that the Japanese yen will appreciate versus the Swiss franc in the period ahead. In actual fact what the value-based investor is saying to himself is that the Japanese yen *should* appreciate versus the Swiss franc in the period ahead and it is my contention that this

market will be covering their shorts very quickly, buying the market back and creating demand. People who were watching the market with a thought of buying, but who were not really convinced, will be saying to themselves, “Well, I was thinking of buying it and now the market is going up. I had better get in now before I miss the boat.” And people who were not really watching the market but had a vague interest in it will be putting down their coffee cups and sitting up to take notice of what is happening. Some of them will be thinking of selling the market and some of them will be thinking that buying the market might be a good idea because it is starting to go up.

It’s like a train leaving the station. Some people are on board before the scheduled departure time but there is always a rush of people to get on board as the train closes its doors and moves off.

Some people I have come across say that volatility breakout strategies do not work, then I discover that they are using historical volatility as their main measurement tool. As mentioned in Chapter 3, historical volatility can be going down when the market is making a trending move and going up when the market is range trading. It makes much more sense to use a measure of volatility such as standard deviation based measures when trying to find “volatility breakouts”. The volatility breakout strategy works and some of the most successful trading systems of all time have been based, in one form or other, on the concept of price contraction leading to price expansion.

VALUE INVESTING

There are actually two ways of looking at value investing as a style. One is to look at it from a point of view, prevalent in stock markets, that investing in a stock with a high dividend yield represents safety and “value”. If we were to think of value in this sense, then it is actually more like a carry strategy where we are investing high-yielding assets, and so in this sense it can be thought of as benefiting from a low-volatility environment.

However, implicit in this idea of value is that the stock with a high dividend yield is, in some way, undervalued and that, by investing in it, we expect that undervaluation to correct itself. Therefore, in this sense value investing is actually a long volatility strategy because it relies upon a market that is deemed to be over- or undervalued, moving back to the level where it can be considered fair value. The investor who is buying “undervalued” stocks, bonds, commodities or currencies needs those assets to move up in order for his analysis to be fulfilled and when he sells “over-valued” assets he needs those assets to fall. Whether they rise or fall on an absolute or relative basis is immaterial. It is still necessary for them to rise or fall (see section on Relativity and Trending Behaviour, p. 343).

There are many different measures of value depending on which markets you trade or analyse, but the one common theme is that they tend to be longer term in nature, and it is this aspect of being longer term in nature that causes them to be bucketed in

very long time indeed and, as such, that has to be taken into account when one is considering buying “cheap” and selling “expensive” assets. As the economist John Maynard Keynes, who probably understood the investment markets better than any of his peers, famously said “the markets can remain irrational longer than you can remain solvent”. However, although the market may take time to return to some measure of what is considered to be fair value, when it does it usually does so quite strongly and in a trending manner.

Thus, successful value investing relies on the market or markets in question giving up on their respective levels of under- or overvaluation and returning towards equilibrium. In actual fact, in the currency markets (and most markets), the move back towards equilibrium value does not end at equilibrium and there is usually an “overshoot” involved, meaning that the market overshoots the equilibrium level and continues on the other way towards the opposite valuation levels. This simply makes sense if we consider the fact that markets are always in a state of flux and change, so that the idea of a market hanging around at what a model claims to be “fair value” is somewhat unrealistic. Yes, a market can hang around at what a model claims to be fair value, but only if the amount of *the market* being supplied is equal to the amount of the market being demanded. This is the true measure of what is value or not in my opinion. If the amount of the market being supplied is equal to the amount of the market being demanded, then, in my view, the market crowd is telling us that this level is, at this particular moment of time, fair value. This, to me, is a more realistic measure of what fair value actually represents. If fair value represents anything, should it not represent where the *market itself* thinks fair value should be rather than where a model or series of models calculate fair value to be? Fair value is a dynamic concept that moves around as the world changes from day to day, and while fair value models try to capture this by the changing values of the variables that are input into them, surely it is the market’s *perception* of fair value that is the most powerful force in deciding where the market will trade. Thus we have periods where supply and demand are similar and therefore the market is in some sort of equilibrium. This is fair value in my book and equates to a trading range in price action.

The overshooting of fair value, therefore, should not be a surprise to anyone who believes that it is the market crowd that is the most important (in fact the only) determinant of market prices and that when a crowd gets into a trend mentality it only stops when supply and demand become balanced again and not when the market price reaches a level at which a model says it should be. Let’s think of an example of a market that has been very undervalued and is finally getting rid of that undervaluation by moving upwards. Is it realistic for us to believe that market participants will think that when the market gets to what is considered by a model to be fair value that they should stop buying or selling their holdings? For sure, some people who were long from the undervaluation levels may well sell their long positions for the reason that they think the market is now at fair value, but will they do so automatically? Will they not be caught up in the trend of the market and buy more? Or will others, who might not have bought the market from the lows and held off buying in

the hope of a pullback, not now buy because the market is at fair value? All of this is possible but what is also very possible is that the psychology of the market drives the price well beyond what is considered to be fair value. A strong trend in motion is a wonder to behold because the psychology of the market crowd is at its most robust when the surge in supply (downtrend) or demand (uptrend) is just so great that the mentality of ever lower or ever higher prices feeds on itself. Does this mean that the market is irrational? Probably. Should that matter to us as investors? Not at all. Why should it? Why should we as investors worry that the market is not rational and not pricing everything at where so-called fair value is? For goodness sake, if everything was priced at “fair value” all the time then markets would cease to exist! Irrationality is necessary for markets to exist and it is what makes markets what they are.

FUNDAMENTAL MODELS

A fundamental model is generally a long volatility style because it again relies on trends appearing in the markets in order to make money, and when we define higher volatility as range expansion, that is exactly what is needed for a fundamental model to add value.

When we talk about a model or system, what we are referring to is a process that happens to be very mechanical and systematic. Most discretionary fund managers, however, will dismiss “models and systems” because they feel that they are too rigid and not nimble enough to avoid potential blow-up situations. They feel that the human brain will be able to process information and come to a better more timely decision than a computer model or system. The ironic aspect to this is that most discretionary fund managers make a big deal in their marketing literature, and when they are “pitching” for assets to manage, make a lot of the fact that they have a “robust and disciplined process” in order to “maintain consistency” when they are making investment decisions. Now, forgive me for being stupid here but is a computer model or system not going to be more robust, disciplined and consistent than a human being or, perish the thought, a *committee* of human beings making decisions? I would say it has to be, just by the fact that the human involvement is taken away from the conclusions of the model. Certainly, the humans input the parameters, but the output is then consistent and disciplined. Nevertheless, the point I want to make is that a model or system is merely a process of decision rules set down in a hard computerised format rather than a set of guidelines that human beings or investment committees generally try to stick to.

Therefore, when discussing fundamental models or systems I am lumping in the discretionary fund managers or individual traders who follow a fundamental “process” when making investment or trading decisions. Whether it is a computerised model or a discretionary process the fundamental analysis of stocks, bonds, commodities or currencies aims to make money, generally, by anticipating movements in the markets, and that means anticipating trends.

is a very, very different thing from saying that the yen will *probably* appreciate versus the franc. The semantics here refer to the fact that value-based investors tend to try to *impose their will* on the market instead of *listening* to what the market is saying. Their conclusion is that the yen is undervalued and that the franc is overvalued, so what is undervalued *should* rise and what is overvalued *should* fall. This is what they have been taught in business school: what is undervalued is cheap and should rise and what is overvalued is expensive and should fall. However, whether things *should* happen is pretty irrelevant if the market is telling us something different (i.e. if the yen, is still in a downtrend and the franc is still in an uptrend, then that is the market's way of telling us that what should happen is a low probability affair!). Nevertheless, value-based investors generally want to buy what is cheap and sell what is expensive in the expectation that somehow (*hopefully!*) the market will realise this anomaly and result in a trending move up in the cheap thing and down in the expensive thing.

In this way value-based investing should be considered to be a long volatility strategy because the investor is expecting a sizeable move from what he considers to be extreme levels. If we take the example of Purchasing Power Parity in the currency markets, we can see that when a currency is extremely under- or overvalued it does generally move back to fair value in a trend-like fashion. Eventually! Figure 23.10 shows the PPP for the GBPUSD exchange rate, and we can see that there can be persistent deviations between what is estimated to be fair value and the actual market price. I do not want to labour the point with respect to value-based investing, but the fact is that a market can remain under- or overvalued for a very,




Figure 23.10 GBPUSD Purchasing Power Parity

Fundamental models and processes exist in all asset classes, and a few examples of the type of analyses used as inputs for these processes are shown below. Please note that many, many more factors will be input into a model, reflecting the fundamentals of these markets. This list is just to provide a very brief flavour.

Currency fundamentals

1. The balance of the external position of the economy. The balance of payments in and out of the currency.
2. The monetary policy of the central bank (relative to other central banks).
3. Gross domestic product growth relative to other countries.
4. Inflation trends relative to other countries.

Stock fundamentals

1. Price/Earnings (P/E) ratio
2. Price to book ratio
3. Price to cash flow ratio
4. Return on equity

Interest rate/bond fundamentals

1. The monetary policy of the central bank of the currency being analysed.
2. The fiscal policy of the government.
3. Current and forecast inflation.

Commodity fundamentals

1. Analysis of physical supply and demand of the commodity itself (note that this is different from an analysis of supply and demand of the commodity future or instrument, which is what a technical analyst does).
2. Analysis of weather patterns and forecasts, which may influence crop-growing conditions.

Fundamental models are therefore very similar to value models or processes because they seek to determine whether a market or security is fairly valued or not in the sense of the dynamic fundamental changes that are occurring. If the model or analyst identifies changes to some fundamental element within the process, then this might lead to an output or view being taken that reflects a bullish or bearish outlook for the market or security. Thus, the fundamental model or analyst then forecasts that the market or security's price will move to reflect this change. This move, if it is generally anticipated, will occur in a trend-like fashion and so, for this reason, fundamental models and processes can be thought of as long volatility investment styles.

SUMMARY

We have considered how certain investment styles can be helped or hindered by certain volatility regimes and how the success of all investment strategies are affected by the volatility trading regime under which they are used. In the next section we will consider how trading regime analysis can help in another major aspect of the investment world.

SECTOR ROTATION AND RELATIVE VALUE

Relativity and Trending Behaviour

A number of years ago I attended a conference sponsored by the UK Society of Technical Analysts and at that conference I heard one of the speakers allude to something that I found to be very interesting. The speaker's name was Nick Glydon, a well-known and respected technical analyst, and he alluded to the fact that although a lot of individual stocks tend to exhibit a great deal of range-trading behaviour on an *absolute* basis, on a relative basis they do trend very well. That is, the absolute value of, for example, IBM shares might seem as if it trades around in large sideways trends but if we looked at IBM relative to its sector or the overall market, there would be strong up and down trends in evidence. He made the observation that another market that tends to exhibit strong trend-like behaviour is the currency market, and he made the connection that, of course, the price of a currency is merely one currency *relative to* another. This realisation is not new to technical analysts in the stock markets and many base their work on examining relative trends of stocks and sectors.

Since that conference, however, I have been investigating more closely whether *relative* asset pricing does indeed exhibit trending behaviour, and the evidence does seem to suggest that there is a basis to this observation. However, the question has perplexed me as to why something should exhibit greater trending behaviour in a relative environment than in an absolute environment. One reason might be that the long-only institutional fund managers who dominate the global investment markets act in a relative manner because their mandates and performance measurement are based in relative terms. This is undoubtedly how these large fund managers will think about, and act in, the market, but why would that in and of itself result in trends that can be extended for long periods of time? Just because a group of people decide to act in the same way at what would probably be around the same time, does not mean that there will be a persistent trend in that market. Even to think that it would, in my opinion, displays an ignorance of how markets actually work.

One phrase that has always annoyed me when people talk about the markets, and particularly in relation to technical analysis, is that "it is a self-fulfilling prophecy". What people mean when they say this is that, for example, if everyone in the market is

orders and not many sell orders at or just above 200p. The technical community buy the market and the price spikes up a little bit. However, what happens then is that the market price drifts back down towards, and in fact dips below, the 200p level. What do our self-fulfilling prophecy people think about that? If all the technical community bought it between 200p and 220p, then the drop in price is hardly benefiting them.

In fact, it is my contention that it is this *drop* in the price that could best be considered the self-fulfilling prophecy because it fits with how any market would behave given the prevailing supply and demand mechanics. What is happening here is that the technical community (for the sake of bucketing them as such to satisfy the self-fulfilling contention) have all bought Tesco as it broke up past the 200p resistance and this increase in demand has resulted in a price spike to around 220p. But then what happens? All the buying has been done and there are no more buyers. There is then a slackening of demand from the abnormal levels seen from 200p to 220p and, assuming a normal level of supply, the market can only do one thing. It has to drop.

The market does indeed drop but finds support around the 200p level and in the subsequent months takes off in a sustained trend towards the 290p level. Who is behind this buying to take it on such an upward trend? It cannot be the technical community because they all bought it on the break above 200p! So if the technical community are already long the market and still the market goes in the direction that the technical analysis points to, then there must be other demand forces at work to propel the market higher. If that is the case, then how can technical analysis be considered to be a self-fulfilling prophecy?

This philosophy is the same when it comes to the case of relative investing behaviour by large fund managers and trending behaviour. These fund managers *do* buy and sell stocks and sectors relative to one another at around the same times, but there must be many more supply and demand forces at work to create the extended trends seen in relative stock and sectors.

The stock market discounts the future – that is, it anticipates the future. That is its job. It exists to act as a barometer of the current and, crucially, expected news about the economy. This is why, when all the news appears to be really good the stock market is falling and when all the news is really bad the stock market is rising. In this sense it could be said that the stock market leads the economy. I have had this argument countless times with my economist friends who persist in focusing on the backward looking data in an effort to “predict” the stock market. The fact of the matter is that the stock market and the economy move up and down in a fairly coincident manner, and a quick check of the charts will confirm that.

Figure 23.11 shows the annual changes in the Dow Jones Industrials versus the annual changes in US Gross Domestic Product. It is clear to see that the stock market and the economy move broadly coincidentally, with the stock market leading the economy more often than not.

If the stock market does not, in some sense, lead the economy then why does the Organisation for Economic Co-operation and Development (the OECD) include it in

watching a particular support or resistance level and act by buying or selling at those levels, then it has the effect of being a self-fulfilling prophecy in that the price will go the way that the analysis suggests purely because of the fact that everyone is buying or selling at those levels. In one sense it could be argued that it will be self-fulfilling if a section of the market is all concentrating on one particular price level and decide to act there, but the questions I always ask people who charge technical analysis with being a self-fulfilling prophecy are: Once the technical guys have bought or sold, then what happens? Does the market just do nothing or does it continue in the direction of the technical analysis? It might and it might not, but if it *does* continue in the direction of the technical analysis then who is driving the price to those levels? It cannot be the technical community because they have already put their positions on according to the self-fulfilling prophecy hypothesis. If the market continues a price move for an extended period of time after an initial support or resistance zone is broken, how exactly can technical analysis be considered to be self-fulfilling? As an example, we will look at Figure 4.8, which is reproduced here.




Figure 4.8 (reproduced) Tesco weekly – classic rectangle

This figure shows the share price of Tesco and the fact that it was trading in a classic rectangular pattern in 1998 through to 2000. Now, if the self-fulfilling mob are to be believed, every technical analyst who was watching this stock would have identified the breakout level around 200p and would have bought the stock as it broke above that level. Let's assume that this is true and that all the technical community bought Tesco as it moved up past 200p resulting in the price spike up to around 220p almost immediately after the break. In this sense it could be said that there is a self-fulfilling aspect to the price *spike* because there would have been a cluster of buy




Figure 23.12 Leading Indicator – Dow Jones industrials

of that sector will turn out to be higher than average. Generally, it turns out that the pricing of stocks and stock sectors do a good job in anticipating future prospects for the underlying company or economic sector.

Another reason why relative values exhibit a high degree of trend-like behaviour could be the fact that when we as human beings think of assets we tend to think of them from a relative value point of view. Let's take a house as an example. At the time of writing, most of us have grown up in the post-1930s inflationary decades which, as result has caused us to think (in most developed societies at least) of inflation as something that is set in stone in economics. We have been conditioned to think that inflation is an economic fact of life and the only problem that central banks have is in controlling the amount of inflation in the system and not letting it get too high. In fact a lot of people believe that the central bank system relies on inflation and inflationary expectations to keep the fiat money system afloat, and this is one reason why the thought of a sustained deflationary period is probably the scariest prospect for central bankers. Deflationary periods tend to coincide with depression or extended recessions in economies and when that happens, the faith in governments (and as a corollary the government backed "fiat" money) could be damaged beyond repair as people run to "hard" currency assets such as gold. This is what happened in the early 1930s in the USA before the government there made it illegal for individuals to own gold and effectively confiscated (some would say stole) their gold holdings! Do you think governments will stop short of unimaginable, draconian measures when they have an agenda in a crisis? Think again.

In 2007 the vast majority of market and economic analysts thought that deflation and depression will never happen again in the west because the central banks will

come to the rescue of the situation, and if all else fails, as the then Chairman of the Federal Reserve Ben Bernanke said, they will just “throw money from helicopters”. This sounds great but it did not work in Japan in the 1990s and early 2000s when they were going through a deflationary depression and land prices were falling year upon year. The “Keynesian liquidity trap” is an economic theory posited by John Maynard Keynes which basically posits that there may come a time when monetary policy has no effect on prices because the confidence of the people has been affected so greatly that they keep their assets in short-term money deposits or even “under the bed” rather than invest in real or other financial assets. This seemed to be what happened in Japan and yet almost everyone, apart from a few notable exceptions, thought in 2007 that a deflationary depression could never ever happen again in the western economies.

The definition of inflation is something that economists argue over. Is it the percentage rise of a basket of goods and services or is it the rate of change of the money supply? Whatever definition is used, however, the final analysis of whether prices will be rising or falling in an economy comes down to confidence and expectations about the future. Once people start to think prices will fall, we get the situation that prevailed in Japan for many years where consumers simply did not spend their earnings because they thought goods and services would be cheaper in the future. By the same token, once people start to think that prices will rise, they will be more inclined to go out and spend their cash because they think things will be more expensive in the future. These expectations are *the* most powerful dynamic when it comes to inflation or deflation, because once they set in it becomes very difficult to break them.

Therefore, over the last few decades in the western world we have become used to inflation and we expect it. Our investment psyche is ingrained in the notion that prices in general will rise forever. Of course there will be exceptions, such as personal computers, but, in general, we have expected prices to rise over the course of time, which brings us back to our house example. House price inflation has been rampant in the western economies for a number of years and, indeed, has been on the rise steadily since the end of the depression in the 1930s. Our expectations that the value of our house will rise over the course of time has reached such proportions that anyone who thinks that house prices might have a *sustained* (multi-decade) period of deflation is considered a crank! Even people who do think that house prices might not rise every year think that there could only ever be a slight easing in house prices, and not a sustained drop.

And this is the point about relativity. In some ways we have become conditioned to think of house prices as a relative value play and a lot of this is down to our inherent belief in inflation. What I mean is that the majority of us don't think house prices will collapse because it is not their *absolute* value that is high, it is their value *relative* to other things, like salaries and earnings, that is out of whack. What needs to happen is that house prices need to go sideways for a period to let their affordability come down in relation to salaries and earnings. Thus we think of assets like houses

in relative terms rather than absolute terms. The deflationists among the readers will notice, of course, that the argument of house prices only going sideways in order to bring their relative value down relies on the fact that salaries and earnings will continue to go higher – a variable that may or may not occur but, at the time of writing, has only higher expectations.

Thus, we tend to think of stuff in relative terms. I can afford the house (i.e. the mortgage payments) because my earnings are such and such. I won't go on that holiday this year because this other one is nearly the same and less expensive. I'll buy stocks now because they are cheap relative to bonds. I will sell my bank sector stock holdings because the retailers look cheap in comparison. This type of thinking is a major factor behind relative trends in asset prices, and if there are trends in *relative* asset prices then trading regime analysis techniques should be able to add value to timing when those relative prices may be in line for a major trend, or when they are more likely to stay in a range-trading regime.

Figure 23.13 shows the chart of the S&P 500 US Homebuilders sector or industry group relative to the S&P 500 Index from late 2004 until early 2008. This chart is using the weekly time fractal and the dotted lines are the 5- and 40-week moving averages of the ratio. The lower chart shows the Bollinger Band Width.




Figure 23.13 US Homebuilders relative weekly and Bollinger Band Width

There are two things to notice from this chart. Firstly, relative trends can last a long time, as we can see from the fact that the weekly moving averages turned down

for the Homebuilders in early 2006 when the ratio was at 0.88 and, although the averages crossed back higher briefly in February 2007, this measure of the long-term trend would have done a very good job of keeping a fund manager short or underweight the Homebuilders during a time of significant underperformance. The second thing to notice is that the measure of trading regime, in this case the Bollinger Band Width, does a good job in identifying periods of trends and consolidation.

The Homebuilders, for instance, entered 2006 in a relatively quiet relationship with the S&P 500, meaning that the Bollinger Band Width was low from a historical perspective. For the trading regime analyst this is a sign that, although things are quiet and range trading is the order of the day, the probabilities favour a reversion in volatility and a return to more of a trending trading regime. We can see from the chart that the next six months produced a period of significant underperformance by the US Homebuilders sector relative to the S&P 500 Index, with the strong downtrend producing a move up in the band width until it started to peak around August 2006. There then followed a long period of consolidation in the relative performance of the Homebuilders during which the band width moved back down to those historical low levels again. When the ratio finally broke out of this range around March 2007, there followed a big collapse in Homebuilders relative performance, with the ratio reaching a low of 0.21 in January 2008.

It is possible, therefore, to use the same analysis of trading regime on relative time series charts as can be used on time series charts of each absolute market. The analysis is no different. If one market relative to another is going through a period of being very steady (range trading), it increases the odds that the relationship will experience a trending regime in the subsequent period. Volatility with relative charts or spreads is mean reverting too, and so the fund manager whose fund is performed on a relative basis, or the pairs (spread) trader who takes advantage of the changing relationship between two securities, can benefit from an analysis of the volatility, and not just the actual level, of the relationship.

If a fund manager has an idea to go overweight or underweight a particular sector then, in my opinion, he would be fairly reckless, firstly, not to look at the underlying trend of the relative performance of that sector and, secondly, not to take account of the volatility of the relationship.

Figure 23.14 shows the line chart of the US Department Stores sector versus the S&P 500 Index using the monthly time fractal from 1993 until 2008 and the associated moving average envelope analysis as discussed in Chapter 9. In this example we are using a 12-month moving average with upper and lower bands placed at 5% from that average. The US Department Stores sector in 2008 consisted of such companies as Kohl's Corporation, Macy's Inc. and JC Penney, some of the most important names in the retail business and, as such, their share price performance is a good barometer of the underlying health of the US consumer. As discussed in Chapter 9, the trading regime logic of moving average envelopes is to identify when demand or supply is strong enough to push the price above or below a certain band width calculated as a percentage around a moving average. The analysis is used to




Figure 23.14 US Department Stores relative monthly and MA Envelopes

identify potential volatility breakouts and trends. On the other hand, when the price is contained within the bands, the market is deemed to be trading in a range-trading regime. We can see from the chart that had an equity fund manager used this analysis in his process to decide whether to be overweight, underweight or neutral the Department Stores sector of the US stock market, he would have been helped to a great extent. In December 2004, for instance, the ratio emerged outside the bottom of the moving average envelope (the bottom band being 5% away from the 12-month moving average) and this was a signal that a potential trending move down in the performance of the department stores relative to the S&P 500 was under way. The sector then underperformed the S&P for another 14 months until February 1996 when the sector relative moved back into the moving average envelope bands.

We can see that there are times when the ratio moves out of, or back into, the bands only to move back in or back out of the bands very quickly. This is an unavoidable issue when dealing with the markets in general and with moving average techniques in particular. This is not a hard science, and moving average envelopes used in this way for trading regime analysis are not supposed to be a trading system with hard rules. Instead they are to be used as a guide in conjunction with other analyses because all the envelope analysis is doing, in common with *all* other analyses, is giving a (subjective) probability reading of something happening. In common with other trend-following types of analyses when the price (or ratio in this example) moves out of the bands and into an anticipated trending regime, the times that it does stick in that trending regime more than make up for the times that it gives a false signal. In other words, when the signal is right it tends to signal a big move in the

price or ratio being analysed but when it is wrong it tends to reverse that signal very quickly. In this way it acts as a normal trend-following type indicator where the number of “winners” are lower than the number of “losers” but the average “win” is much bigger than the average “loss”. The analyst, therefore, has to be pragmatic when using moving average envelopes to determine a trading regime and should expect a number of false signals. This is why I like to look at the closing price only when looking at moving average envelopes, something that the Turtle type non-continuous trend-following system did too. If we look at the closing price only, we can get rid of a lot of the noise that is associated on intra-bar trading. Intra-bar trading refers to the bar chart where the open, high, low and closing prices are shown and so there will always be a range of prices on each bar. If we use bar charts on moving average envelope analysis – in the sense that we are using them to determine trending and ranging regimes on account of whether the price is inside or outside the band – there will be a great deal of overlap inside and outside the bands. If we just used the closing price, however, we can minimise this noise and only focus on the closing price that, as alluded to earlier, is regarded as the most important price at least from a psychological viewpoint. The closing price is the price that, at least on a daily and weekly time fractal basis, everyone goes home to bed and thinks about. They either feel good or bad about it depending on which way they are positioned, and these human feelings have a direct effect on how they will act in the market the next day or week.

Another benefit of just using the closing price, at least when using daily or longer data, is that the closing price, being the one to which firms predominantly use to mark to market their investments, will generally always be correct. What I mean is that it will be a true reflection of where the market closed. The problem with using high and low data is that there may sometimes be discrepancies as to whether the market actually traded there or not. This happens a lot in the foreign exchange market but can happen any time there is a “fast” market, be it in stocks, bonds or commodities. The data on our quote screens might say that the low in, for example, the EURUSD exchange rate on a particular day was 1.5520, but if the market had spiked down perhaps after an economic data release – that is, rather than having traded at 1.5520 it actually traded at 1.5525 (the 1.5520 just being a “quote” that was never traded on) – it can make a difference to someone’s trading system or process. It may only be a matter of fractions, but if it is not the true price it does not truly reflect the true psychology of the market. For technical analysts, therefore, whose *raison d’être* is to analyse the market price action, the integrity of the price data must be paramount and the integrity of the closing price, at least on a daily basis, is generally accurate because it is used for valuation purposes.

I don’t want to labour the point, but one of the things that really “gets my goat” is when the fundamentalists tell me that technical analysis is fine but it is only really for short-term analysis. If anyone tries to tell you this then sit them down (even if you have to physically restrain them!) and show them some charts using the monthly time fractal that will have obvious trends and price patterns. A price is a price is a price. It doesn’t matter whether it is the closing price every minute or the closing

price every month. Human psychology is affected by that price, and it just depends on which group of people look at which price. For most people the one-minute price is totally irrelevant, but for the majority of investors the one-month price is extremely relevant. Joe Sixpack's mutual fund statement will come in each month, whereupon he will look at it and, if it has gone up in value he will be happy but if it has gone down in value he will frown. So what does Joe do? If it has gone up he might think "I should invest some more" and if it has gone down he might think "I am worried about this, maybe I should get my money out". This is natural behaviour and is allied to the theory of positive shaped demand curves in financial markets, as mentioned earlier. The point is that the human psychology is affected by price data no matter how frequently it comes, and if human psychology is affected by low-frequency price data, then actions in the market that are a direct result of that psychological dissemination of the data will also be a direct result of the low-frequency price data.

Therefore, technical analysis and trading regime analysis are not only relevant for the longer term investor but are, in my opinion, of utmost importance because price trends using longer term, low-frequency data will, by their very nature, last a long time and being on the correct side of these trends will be vital if someone is to perform well in the markets. For longer term equity fund managers, whose job it is to perform on a relative basis, being on the right side of big trends in a sector performance relative to the overall market can make a huge difference.

Figure 23.15 shows the relative performance of the US Pharmaceutical sector versus the S&P 500 Index using the monthly time fractal from 1990 to 2008. This sector comprises well-known blue-chip companies such as Johnson & Johnson, Pfizer Inc., Abbott Labs and Merck. At the time of writing, this sector (or subgroup) of the S&P




Figure 23.15 US Pharmas relative monthly and MA Envelopes

500 had the second highest weighting in the index (at 6.225%), meaning that the S&P 500 Index's performance would be greatly affected by how this sector was performing. It also means that for a relative fund manager the decision of whether to underweight, overweight or be neutral this sector is extremely important in terms of risk (or tracking error as it is known in portfolio management) and return.

The chart also shows the moving average envelopes using a 12-month moving average and 3% bands either side of that, and we can see instantly that this ratio (the Pharmaceutical sector divided by the S&P 500) can move in extended trends but can also stay quite stable for a time. Using the usual logic of applying the moving average envelopes for trading regime purposes (i.e. as a non-continuous trend-following indicator) we can see that this would have been indicating to a fund manager that he should have been underweight the "Pharmas" from April 1992 until early 1994, when the sector underperformed the S&P 500 Index quite badly. After a brief period of the ratio trading in between the bands when the indication to our relative fund manager would have been to be neutral the sector, the Pharmas then began a long period of outperformance from August 1994 that lasted, give or take a couple of very brief dips back into the bands, until 1999. Of course, due to the discretionary nature of using this analysis as an indicator rather than a hard system, all the value of the outperformance may not have been realised. That would depend on how the fund manager incorporated the indicator into his overall investment process. However, this was a strong multi-year trend and the moving average envelopes gave a good indication of when it would be prudent to have had an overweight position in the Pharmaceutical sector in the US.

The sector did not waste too much time in a range-trading environment when it went back into the bands, and instead emerged on the other side in April 1999 from where the Pharmas underperformed badly once again for another year, before turning back up from 2000 to 2002. From January 2002 until July 2003 it stayed mostly within the bands, where the indication would be to be neutrally weighted, but in the summer of 2003 the Pharmaceutical sector began another long period of underperformance that lasted until around June 2006. I know people reading this will be looking at the data on the chart and saying "look, it went into the bands there, indicating a potential range-trading period and popped straight back out" or "it came out of the bands, indicating a potential downtrending period only to go straight back in again". I cannot reiterate enough that although something like moving average envelopes could be used as a non-continuous trading system with hard rules, that is not how I prefer to use the analysis. Instead, I prefer to let the analysis guide me into how I should *probably* be positioned and combine this with my other analysis before coming to an investment decision. Again, a multi-time fractal approach helps in this regard with an analysis of this monthly chart being used as "the guide" and the actual investment decision being based on the weekly or daily chart.

As an example, let's look at the false signal of a potential major downtrend beginning in April 2002. The ratio closed the month of April 2002 outside the bands only to re-enter the bands three months later in July 2002. This turned out to be a false

then, even though the monthly analysis was still suggesting a potential downtrend, the fund manager would not be underweight the Pharmaceutical sector but instead be neutral. At the end of April the ratio was 0.3510 and this, therefore, is the initial ratio level at which the fund manager might have gone underweight. On 26 July, however, the Pharmas came roaring back and the moving averages turned back up. At this point the fund manager would, according to his process, move from an underweight position to a neutral position on the sector. On that day the ratio closed at 0.3530 so the fund manager would have emerged roughly flat, or with a tiny loss. Either way he would not have held on to an underweight position as the market continued to rally over the next few weeks and, in this multi-time fractal way, the risk management of a market position can be run effectively.

Relative Interest Rate Spreads

The concept of volatility analysis can also be applied to interest rate spreads. Many traders and fund managers try to take advantage of the relative movement in interest rates or bond yields between two different countries, between different bond sectors (like government versus corporate bonds), between different maturities along the yield curve or a combination of all of these. Again, the same reasoning applies. If a relationship is very stable, the chances are that instability is coming next and if the relationship has just experienced a big rise in instability (a trend) the chances are that a period of quiet (or at least trend exhaustion) will appear next. This is important for traders who want to know about the odds of a potential trading regime before putting a trade on because the same trading methodologies apply to these spread strategies as to each individual market. If the odds favour a trending regime, then our spread trader will know that it might not pay to “fade” (to look for a reversal) a breakout of recent ranges in the spread. Equally, if the spread has just experienced a large up-tick in volatility, our trader will know that the odds favour an exhaustion of that trend at some point and, if not a trend like reversal, at least a diminution of the directional volatility of the preceding trend (if it was an uptrend then odds favour, at least, a slowing or stalling in upside volatility) which could manifest itself in a range-trading regime.

Figure 23.17 shows the interest rate spread between the US and European 10-year government bond yields in the weekly time fractal from 2001 until 2008. This is simply the US government bond yield minus the European government bond yield; this means that when the spread is going higher the US bond market will be underperforming the European bond market in price terms (because the US yield is either going higher quicker or lower slower than the European bond yield) and vice versa. The lower chart shows the vanilla 20-period standard deviation of the spread and we can see that periods of ranging behaviour are coincident with falling and/or low standard deviation, whereas periods of trending behaviour are coincident with increasing standard deviation. Some of the periods of low standard deviation give

signal of the start of a potential major downtrend. However, would the fund manager have actually gone underweight the Pharmas as a result of the monthly close outside the bands? Quite possibly, but if the fund manager was using a multi-time fractal approach, there is a good chance that he would not have taken the investment decision to go underweight.



Figure 23.16 US Pharmas relative daily

Figure 23.16 shows the US Pharmaceutical relative chart but this time in the daily time fractal (daily closing prices) and two moving averages, the 5-day period and the 40-day period. Let us suppose that our fund manager was using an investment process whereby the monthly chart acted as a guide, but the analysis of the daily time fractal trends was where the investment decision was made. In other words, he would not act in the market unless the daily time fractal gave a trend-following signal that was consistent with the trading regime analysis as presented in the monthly time fractal (in this case a downtrend). Therefore, he would only go underweight the Pharmaceutical sector when the 5-day moving average turned down below the 40-day moving average from above. We can see from the chart that in April 2002 the 5-day moving average was already below the 40-day moving average, indicating that a downtrend was in force. Would the fund manager have initially gone underweight because of this? Perhaps, and probably it would be prudent to at least have taken a small position to hedge against the possibility that the monthly signal resulted in a complete multi-month collapse in the Pharma relative without any retracements within that trend. However, the fund manager's process is that he will only be underweight the Pharmas as long as the monthly moving average envelopes indicate a downtrend *and* that the daily moving averages are indicating that a downtrend is in force. What this means is that were the daily moving averages to turn up (the 5-day above the 40-day)




Figure 23.17 US–EU 10-year Yield Spread weekly and Standard deviation

rise to sharp trends (volatility breakouts) and some give rise to longer trends (such as after the long consolidation between October 2005 and September 2006).

What, however, should a trader or fund manager do with this information? Again, it is all about risk management and probability. If the band width is very low, the odds favour a volatility breakout of the spread; and if it is very high, the odds favour an exhaustion in the recent trend. Incorporating this volatility analysis into their process can add an extra layer of value to the interest rate trader.

24

Trading Regime Analysis for Economists and Fundamentalists

“Stock prices have reached what looks like a permanently high plateau. I do not feel there will be soon if ever a 50 or 60 point break from present levels, such as (bears) have predicted. I expect to see the stock market a good deal higher within a few months.”

Irving Fisher, Professor of Economics (17 October 1929)

There is no reason why some concepts of technical analysis cannot be used on some economic and fundamental time series. If we accept the hypothesis that economics is a study of human behaviour in the sense that it captures the outcomes of human choices at the societal level, then we can accept the fact that trends will probably exist in at least some of this data. I mean, what economist would tell you that trends do not exist in the data? Over the very long run I can certainly accept the fact that most economic time series are mean reverting in nature as the data captures the ebb and flow of the business cycle, but even for something that is mean reverting it has to have a trend in order to get to a point where it is going to mean revert. And what fundamental analyst is going to tell you that there are no trends in a company's earnings or ratios like return on equity? There are trends that exist in these measures because these measures reflect the *change* that is constantly occurring. The point is that these fundamental or economic data series are, in the final, final analysis, driven by human beings making economic decisions, and so the movement in these time series will be ultimately driven by human psychology at both the individual and crowd level. The only difference between the fundamentalists and the technical analysts is that the fundamentalists seek to anticipate change through one set of data and the technical analysts seek to anticipate change through another.

In fact, the concept of trading regime is inherent within orthodox economic analysis because it is the rate of change (be that per month, per quarter or per year) that is important in many, if not all, economic time series. Economists are constantly telling us that the month on month rate of change is $x\%$ or that the quarterly rate of change in GDP is $y\%$. This focus on the rate of change is mainly due to the fact that most economic data, like GDP or production, is expected to be going up over the long term. Extended periods of economic contraction are ruled out by most orthodox economists because the belief is that central banks or other such authorities would have the ability to “turn the economy around” back on to the long-term “trend growth rate”. Up until very recently, when the extent of the credit contraction of 2007 and 2008 has become clear, anyone even considering anything like a depression scenario in the western economies was considered to be eccentric. Time will tell whether the credit contraction will have a lasting effect but at the time of writing there are at least a tiny minority of orthodox economists willing to accept the fact that extended economic contraction cannot and should not be ruled out totally.

Nevertheless, it is the rate of change that is the predominant analysis tool, which most economists take to be the preferred choice. In effect, by doing this economists are actually measuring the trading regime or volatility of the economic time series in question. If we applied the same logic of the rate of change we used to market data, and if an economic time series’ rate of change was close to zero, we would then call that a low volatility range-trading regime. Crucially, however, we would expect the next regime to be a trending regime when the rate of change moved away from zero. On the other hand, if the rate of change was quite high relative to history, we would expect some trend exhaustion to appear and that the next period would be one of range trading or low volatility (the economists call it cooling off). The problem is, in my experience, that most people (indeed even some orthodox economists) tend to forget about these cycles in rates of change and tend to project current rates of change into the future. For example, if retail sales are expanding at a 5% annual rate, then most people will expect this to continue and forget about the fact that the 5% annual rate has been the peak in rate of change in the past when, after that, a period of “cooling off” and a downturn has taken place. This cycle in volatility of economic time series is generally known as the business cycle! Thus, the concept of rate of change is inherent within economic analysis, but what about other trading regime techniques? Let’s look at how another trading regime analysis technique can be applied to an economic time series.

Figure 24.1 shows the line chart of the monthly United States Industrial Production Index from 1988, and moving average envelopes using a 12-month moving average and 1% bands. The Industrial Production Index measures real output and is expressed as a percentage of real output in a base year. For economic time series such as these it is possible to analyse them from a trading regime perspective, and by examining via the lens of something similar to moving average envelopes the economist can

statistics were behind us and that the industrial production in the US was bottoming out? It certainly would have made an economist less pessimistic because it is a recognition that a strong downtrending regime was not persisting. Remember, trading regime or volatility analysis is not supposed to be predictive like the economists try to be. It merely reflects the reality of the volatility of the time series, but in doing this it gives the analyst a probability reading of the current and *anticipated* future trading regime.

Industrial production in the US traded within the bands until March 1992 when it popped out the top band. At this point the economist using this analysis would observe that industrial production was showing strong momentum, strong enough to propel it above the upper band level, and if our economist is an observer of trends, as most are, then he might conclude that the probabilities suggest a continuation of this upward-trending regime. In reality, of course, our economist would not blindly accept the fact that because the time series was showing strong momentum it would continue. Oh no. He would spend most of his time searching for “reasons” and justifications why it would or would not continue rather than just follow the trend. Nevertheless, the strong trending regime did continue higher from this point and, with only the briefest touches of the upper band on three occasions, industrial production in the US remained in a strong uptrending regime throughout the 1990s. In October 2000, however, the momentum of industrial production had weakened to the point that the time series had gone back into the moving average envelopes. To the economist using this trading regime analysis this would have been quite interesting information because it was the first time in eight years that the series had lost enough momentum to put it back into the bands. Does that information alone give the predictive economist any edge? Possibly not, but for something to have had such a long period in a strong trending regime a move out of that regime should be noteworthy. As it turned out this loss of momentum and move back into the bands did not last very long as the industrial production series emerged below the lower band in January 2001. At this point our trading regime economist is presented with a series that is in a strong downtrending regime and therefore, using the logic of the trend follower, will continue to trend down. It did so for the next year until, in March 2002, the industrial production in the US had stabilised enough to get back into the moving average envelopes.

Although, since then, there were a couple of periods when the series could have been said to have been in a strong uptrending regime as it was trading outside the upper band, the vast majority of the time the industrial production series has been contained within the envelopes. For the economist this should be an interesting development because it is essentially saying that the momentum of the industrial production series was much, much stronger in the 1990s than it has been in the 2000s. In other words, the uptrend in the 2000s has been much weaker than the uptrend of the previous decade. Does this mean that US industrial production growth has had an underlying weakness and is therefore set to drop? Quite possibly, and this data will be worth updating in the years to come I am sure.




Figure 24.1 US Industrial Production and MA Envelopes

get a sense of whether the trend is not only up or down, but whether it is strong or weak. Most economists would probably retort that they always do this instinctively when they look at something like a moving average. If they see the moving average going mostly sideways they know it is a weak trend. However, by using something like moving average envelopes they can formalise the process a little more by having more of an objective measure of the volatility (expansion or contraction) of the time series. We can see from the chart that in 1987 and 1988 the United States Industrial Production was in a strong uptrend because it was trending higher above the upper band of the envelope. However, in late 1988 the time series moved back into the bands and this would have indicated to the economist that the strong uptrending regime of industrial production in the US was possibly coming to an end. Interestingly, industrial production then stagnated for the next couple of years, trading within the 1% bands of the 12-month moving average. Was this stagnation a sign that the economy was slowing down and a precursor to the recession of 1990? Perhaps. All economic indicators slow down and weaken before recessions are officially recognised so we can't point to the fact that the industrial production series was in a low volatility range-trading regime since the end of 1988 as a sign that recession was definitely coming. However, we can say that industrial production was no longer in a strong uptrend and therefore that the economy was slowing down. The series popped out of the bottom band in November 1990 but only stayed there for seven months before moving back inside. Was this a sign, in June 1991, that the worst of the official

Therefore, the logic of trading regime analysis and technical analysis in general is actually very well suited for application to some economic time series. Trading regime and technical analysis of time series attempts to identify the underlying psychology behind the data, and if we can accept the fact that a lot of economic time series are merely the graphical manifestation of human beings decisions (just like a market price), the use of such analysis should be valuable to economists.

25

Case Studies

THE PRIVATE SWING TRADER

In the case of a private trader who is trading or investing in the markets in order to enhance his existing income, or in order to produce his only source of income, an appreciation of trading regime analysis can really help results. Let's take the example of John who is a retired engineer and trades the markets in order to enhance his regular pension income. John wants to produce some regular value from the markets but doesn't want to be staring at a screen all day, so he decides that the ideal time fractal to follow is the Minor or daily time fractal. This means that he will take positions based on a technical analysis of daily price data and not be bothered with anything shorter than that. This also means that he will pay attention to volatility analysis using a higher time fractal such as weekly data. As John lives in the UK he decides that he wants to trade the UK stock market because it feels very familiar to him. John feels much more comfortable being a buyer of volatility rather than a seller of volatility, and therefore looks for situations where volatility breakouts seem to have a high probability. It is January 2004 and he notices that the share price of Reuters Group, the news and financial information group, is in a distinct coiling pattern with the weekly Bollinger Band Width falling to historically low levels. Figure 25.1 shows that the share price of Reuters had bottomed out in early 2003 after a steep fall over the previous few years. The market then moved higher in an impulsive Elliott Wave 5 movement to a peak of 741 pence in August before pausing, and since then the market has traded sideways. As a result of this coiling in the price the Bollinger Band Width has collapsed but, not only that, the price has been tracing out a distinct ascending triangle pattern. This suggests that the underlying psychology is still bullish and that, in all probability, the market will eventually break higher above the 741p upper resistance zone in another impulsive move.

John decides that he will go long 5,000 Reuters shares if the market breaks up, and we can see from Figure 25.2 that on 13 January the market moved up and closed

at 753p. Volume had also risen on this day which made John more confident that this break up was a real volatility breakout and not a false dawn. Although his entry strategy is done on a discretionary basis he also likes to use a trailing exit strategy that enables him to follow the trend, and for this he uses a 5- and 20-day moving average cross indicator (the averages are also shown on the chart). John is long 5,000 Reuters shares at 753p and on 15 January 2004 the market explodes higher on massive volume.

This confirms that the volatility breakout has occurred but, as he is not sure how far it will go, he uses the moving averages to trail the exit. We can see that this exit strategy kept John in a long position even though the market traded sideways for a number of days from 16 January before moving higher again. Although the market reached a high of 1213p on 18 February, John stuck to his trend-following exit strategy and eventually exited the position on 10 March at 1023p. This still gave him a return of 41.7% (annualised 371%) and a gross profit on his 5,000 shares of £5,750. Not bad for a few weeks investment.

THE CURRENCY OVERLAY MANAGER

Currency Overlay is a generic term for an investment programme that manages the currency risk that is generated from a portfolio of assets. Those assets could be stocks, bonds, real estate or anything that has an associated currency risk as well as its own specific asset risk when the asset is valued back into the base currency of the fund. The base currency refers to the currency in which the assets are valued, so if I were a US-based investor the base currency of my fund would logically be the US dollar. However, if I have an international portfolio of stock, bond and real estate assets, the currencies in which they are traded will have to be valued back to the base currency of my fund, which is the US dollar. When this currency risk is taken into account, the investment decisions involved in managing the fund become a little more complicated. Not only do I have to worry about the risks associated with investing in stock, bond or real estate markets but I also have to consider the currency risk I am assuming when I invest in these assets. And this currency risk can be very significant if it is not taken into account. For instance, I could invest in the New Zealand real estate market and see the value of those investments soar in local currency terms. However, I must convert the assets back to the base currency of the US dollar, and if the New Zealand dollar has weakened by more than the real estate assets have risen, the fund has actually lost money on the investment overall.

This is very obvious to anyone involved in investments be they professional or non-professional. However, it never fails to astound me just how little people do when it comes to managing their currency risk. This could be due to a number of factors but one of the biggest, at least in the professional sphere, is that people (generally economists and fundamental analysts) tend to think of currencies as tending




Figure 25.1 Thomson Reuters PLC weekly and Bolling Band Width




Figure 25.2 Thomson Reuters daily and Volume

to revert around a long-run average of zero. That is, they may cycle up and cycle down but they always revert back towards a long-run average after some time. In a sense I have to agree with this. Currencies do eventually revert to a long-run average and they do cycle up and down. However, the long-run average does not have to be a straight line and it can be a moving average with a strong long-term direction! If we look at the exchange rate of the US dollar to the Japanese yen as an example, in 1971 the exchange rate was 350 yen to the US dollar and over the next 25 years it fell to less than 100 yen to the US dollar. The average exchange rate over this period was around 200 but since the USDJPY exchange rate fell to around 100 in the 1990s it has stayed there or thereabout with the high being just below 150 in 1998. If long-term reversion to the mean is correct, then should not the USDJPY rate be starting to move up more? No, long-term reversion to a mean should be thought of as a moving average that takes account of structural changes in the world.

Managing currency risk is therefore an extremely important part of the overall investment process and, as more people have realised this, the concept of currency overlay has grown relatively rapidly. There are vast arrays of currency overlay managers who offer their services to asset managers and within asset managers there are specialised areas that do the job internally. When it comes to investment style, the manager or internal desk could have just one style that they offer (for instance, trend following) or they could offer multiple styles of managing risk in the currency markets. This concept of style diversification is gaining popularity as people realise that it is just the same principle as the well-known strategy of diversifying long-only (funds, stocks or bonds that people buy) investments. The underlying reason why people diversify their long-only investment decisions is because they think that if one sector of the stock market is doing badly then having investments in other areas that might be doing well will atone for this. It is the age old concept of not putting all your eggs in one basket and, for some, this is a principle that is set in stone. Like anything else, of course, there is a trade-off and the trade-off is that a diversified portfolio will underperform a concentrated portfolio when the concentrated investments are doing well, but it will outperform the concentrated portfolio when that is doing badly. There is no investment decision that does not involve a trade-off between risk and reward! Nevertheless, in the area of currency overlay style, diversification is a well-used strategy.

As an example, let us assume that Neil, our currency manager, has two styles that he uses to manage currency risk. One is a trend-following style that makes money when the markets trend and the other is a mean-reverting strategy, based on slow stochastics, that makes money when the markets do not trend but instead trade in a range. In order to help to determine how much capital or risk to allocate to each particular style, Neil utilises a mixture of trading regime analyses, and because the trading time frame of the currency overlay programme uses daily data for entry and exit signals, this trading regime analysis is done using weekly and monthly time fractal data. By employing this multi-time fractal approach Neil can take a long-term analysis of volatility conditions and apply that to the actual time frame used in his

day-to-day trading. There are a number of ways that trading regime analysis can be applied in the currency overlay space but they essentially boil down to the decision of whether to analyse the trading regime of each currency pair individually or whether to analyse the trading regime of a currency index. A currency index is merely an index of one currency's performance versus a basket of other currencies. There are many different ways to weight these baskets within the indices but the most popular would be weighting them either by the amount of physical trade a country does with others (trade weighted) or by weighting them by the amount of currency flow that goes through each currency pair in the foreign exchange markets (flow weighted). Even within these methods there can be differences between the weighting criteria, but the differences tend to be relatively minimal. Alternatively, one could construct one's own currency index with perhaps equal weightings for each currency pair. It doesn't really matter how the index is weighted just as long as the logic of finding a feel for the trading regime outlook of the currency in general is not compromised.

The decision of whether to analyse an index of each individual currency pair for trading regime purposes will depend on how static or dynamic the programme is set up. For instance, a static programme would entail allocating a certain amount of risk to each style and each style then trades all currency pairs in the programme when signals are given. For this static programme a trading regime analysis of a currency index would be more appropriate because that would give us a general feel for the likely volatility conditions in the period directly ahead. A more dynamic programme, on the other hand, would look at each currency pair's trading regime individually and decide how much risk to allocate to the style as appropriate. This way requires more analysis, but it also would potentially mean that we could have a situation whereby if all currency pairs were showing one particular regime, then more risk might be applied to that style. In other words, analysing an index-trading regime would mean that the programme might be more conservative in risk allocation.

Let us suppose that Neil has the resources to analyse each currency pair individually and that the base currency of the underlying fund is euros, meaning that the currency pairs that the programme trades are the euro cross rates. The euro exchange rate with the British pound has a relatively high weight in the underlying asset index, and so it assumes quite a degree or importance in the currency overlay programme. It is August 2007 and John is conducting a trading regime analysis of EURGBP.

Figure 25.3 shows the bar chart of the EURGBP exchange rate using the monthly time fractal and with the associated Bollinger Band Width from 1994 until late 2007. The data of the euro prior to its launch in 1999 is synthetically created using the weights of the European legacy currencies within the single currency, and we can see that the EURGBP exchange rate has gone through quite distinct bouts of price contraction or low volatility as well as price expansion or high volatility. The chart shows that since the period of range expansion in 2003 the price went into a very distinct consolidation mode with the monthly standard deviation, as measured by the Bollinger Band Width, dropping relentlessly. Although the band width became extremely low in 2005 (meaning that the chances of range expansion or a trending




Figure 25.3 EURGBP monthly and Bollinger Band Width

regime were increasing) the fact that it was still pointing down would have given Neil some pause for thought before allocating a heavy amount of risk towards trend following and away from the mean-reverting slow stochastics. By early 2006 the band width had dropped to levels not seen since 1996, and in the latter part of that year it finally started to show signs of coming to life again by rallying. At this point in late 2006 John would be quite excited about EURGBP because the long-term volatility of the market was at extremely low levels and was showing signs of coming to life again. In January 2007 the market moved down quite sharply but by the end of the month it had recovered to close at 0.6630. This was an exciting development for Neil because it meant that the potential that all the consolidation in price action from 2003 was a mammoth symmetrical triangle and was turning out to be a high probability. The market had fallen during the month but had recovered strongly to close inside the lower bound of the triangle, and this analysis meant that there might be some more range trading; but eventually, when the break of the triangle came, it would mean a big trending period. The monthly fractal analysis is one factor in Neil's decision-making process but he also looks at the weekly time fractal to give him a more complete feel for the potential trading regime or volatility conditions ahead, and in early 2007 the volatility conditions in that time fractal also began to pick up from multi-month lows. In early 2007 Neil's trading regime analysis was telling him that the really long-term symmetrical triangle was still in force, but that the multi-year extreme lows in volatility, as well as the multi-month lows in volatility, were beginning to show signs of life. At this point Neil had decided on a plan for EURGBP.

He would allocate 66% or two-thirds of the risk budget available in EURGBP to the trend-following programme now, and would hope to increase that to 100% if and when the price broke out of the long-term symmetrical triangular pattern.




Figure 25.4 EURGBP daily and Slow stochastics

Figure 25.4 shows the daily time fractal chart of EURGBP with two moving averages, a 5-day and a 40-day, as well as the slow stochastic of the price in the lower chart. The moving averages represent the very simple continuous trend-following system that Neil uses, which is long when the 5-day is above the 40-day and short when the 5-day is below the 40-day. We can see from the chart that after an initial period of trending in early 2007 the EURGBP exchange rate went back into a range-trading regime for a number of months and this resulted in a bit of whipsaw or chop for the trend-following system. However, in September 2007 the exchange rate burst higher and in the process went up to break the upper bound of the long-term symmetrical triangle. According to his plan that was made at the start of the year, Neil then increased the risk allocation to the trend-following style from 66% to 100% in the expectation that the market would now enter a period of a strong uptrending regime and we can see from the chart that this was indeed the case, with the exchange rate moving higher some 16% over the course of the next seven months.

The slow stochastic on the chart represents the mean-reverting strategy and we can see that when the market was range trading this strategy did well in picking the market's turning points. However, as with all range-trading and mean-reverting

strategies, when the market began to trend strongly it did quite badly in picking turning points. The slow stochastic gives sell signals when the market is in overbought territory but, because the trend is so strong due to the fact that the market is emerging from a multi-year period of extremely low volatility, these sell signals are continually stopped out.

In this example, the trading regime analysis enabled Neil to add a great deal of value to the currency overlay programme by allocating risk to a trend-following style and away from a mean-reverting style when the probabilities of a long-term trading regime shift were high.

THE DURATION OVERLAY MANAGER

Ewan works on the bond desk of fund management organisation and is responsible for the duration overlay of the European government bond market exposure. What this means is that he is responsible for switching the fund's exposure to different parts of the European bond markets as and when he sees fit. A bond's duration refers to the weighted average maturity of the bond's cash flows where the present values of the cash flows serve as the weights. The greater the duration of a bond the greater will be its price volatility, so a 2-year bond will be much less volatile than a 30-year bond. The allocation to the European bond market is invested on an index replication basis, meaning that whichever index the bond fund is performed against, the physical cash is invested into physical bonds in such a way as to exactly replicate the index and therefore the duration of the index. This means that the cash investments and the index will move up and down in the same way. Ewan uses his expertise to overlay derivative investments on this fund by buying or selling European bond market futures and other derivatives on various parts of the yield curve. Sometimes he will use these derivatives to express a view on the shape of the yield curve. If he thinks the yield curve is going to become steeper (meaning that the yield spread between the short end of the yield curve and the long end of the yield curve will, when it is in a normal shape, become larger) he may do something like buy 2-year futures and sell 10-year futures. Alternatively, he can express a view by simply buying or selling one particular part of the yield curve and this will alter the overall duration of the fund relative to the index as well.

It is 14 May 2007 and Ewan has been studying the chart of the 10-year yield on the European government bond.

Figure 25.5 shows the bar chart of the 10-year European government bond yield using the daily time fractal, with the Bollinger Band Width in the lower chart window. Ewan has noticed that the market has rallied sharply from the March low of 3.87% to a high of 4.24% first achieved on 13 April. Since then, however, the market has traded in a distinct range-trading regime with the Bollinger Band Width collapsing to historic lows as a result. However, he also notices that the market has remained capped at the 4.24% level, whereas the reaction lows of the consolidation phase have




Figure 25.5 Euro 10-Year Government Bond Yield daily and Bollinger Band Width

been getting higher. The initial low on 19 April was at 4.12% but since then the lows have been at 4.14% and 4.16%, with the most recent one on 11 May. In this regard the 10-year European government bond yield is tracing out what looks like a classical ascending triangle where there has been quite a strong up move into the pattern and the consolidation phase is showing increasing demand being capped by static supply. According to this analysis the market should be ready to make an impulsive move higher in a trending regime once it breaks the consolidation or range-trading phase, and on this particular day the market is closing above the 4.24% resistance zone. Ewan decides that this provides a good opportunity to reduce the duration of the bond exposure and decides to sell Bund futures contracts in order to benefit from a rise in European 10-year yields and therefore a fall in the price of the futures contract.

As we can see from the chart, the yield did move higher in a high volatility trending fashion, moving from 4.25% on 14 May to a high of 4.70% on 13 June.

THE ASSET ALLOCATOR

Roddy is Chief Investment Officer of a major fund management organisation and his main responsibility is to manage the asset allocation process in deciding whether to underweight or overweight the asset classes in which the funds invest. The broad choices are stocks, bonds, commodities and cash. It is September 2000 and Roddy, our asset allocator, is weighing up the long-term prospects for the four main asset classes of stocks, bonds, commodities and cash. There are many inputs into this

process but one of his main inputs is an analysis of the long-term trading regime because he realises that the performance of asset classes will cycle over the long term and that asset classes that have been relatively quiet over the last few years will tend to come to life, and vice versa. In this regard Roddy is following the volatility that rotates between asset classes over the long term.

First of all he looks at stock markets in which his firm, the fund management arm of a major insurance company, have the majority of their focus and investments. He firstly looks at the Morgan Stanley Capital International (MSCI) Local World Index using the long-term monthly time fractal. Using a graph, as shown in Figure 25.6, he notices that the long-term uptrend global equities have been relentless but that on occasions, such as in 1994 and 1998, the extremely high and falling readings of the Bollinger Band Width acted as a brake on the trend and set levels where the market paused or, in the case of 1998, fell back quite sharply. Roddy notices that, once again, the Bollinger Band Width reached extremely high levels in March 2000 and that since then the global stock market index has been trading in a sideways range-trading regime. However, what really catches his eye in September 2000 is that the Morgan Stanley World Index has traced out what looks to be a classical head and shoulders price pattern. The high band width has led to a consolidation, but that consolidation has resulted in a long-term switch in psychology where the market has tried to make a new high, failed, and has crucially come back down below




Figure 25.6 MSCI World Equity Index monthly and Bollinger Band Width

the neckline of the head and shoulders pattern. In the last few months the long-term psychology and the long-term trading regime of the global stock markets has changed from a strongly bullish trending regime to a neutral range-trading regime, and finally to a potentially bearish trending regime. Roddy factors this information into his overall asset allocation decision and decides that a prudent course of action would be to underweight global stocks at this time. This turned out to a useful decision as the global stock market index crashed by over 40% over the next couple of years.

Now skip forward to July 2003 and Roddy is looking at the Commodities Research Bureau (CRB) Index to help him with a broad understanding of the potential trading regime that commodity markets may face on a long-term outlook. He is looking at the long-term monthly time fractal (as depicted in Figure 25.7) and he notices an interesting development in terms of the market psychology. Since the peak in the CRB in 1980, commodities have been in a very long-term bear market when the index has been making lower highs and, given the new low in 1999, lower lows. It has been a slow sluggish downtrend but a general downtrend nonetheless. However, since that low was made in 1999 at 181 on the index, the market subsequently bounced to 234 in 2000 before going lower again. This is where it starts to get really interesting from a trading regime analysis point of view for Roddy. The CRB index fell again to the 181 level in late 2001 but that level held as support, and the market raced higher from there, moving above the previous high of 234 in January 2003. The market then fell back but, crucially, the old high of 234 is now acting as support




Figure 25.7 CRB Index monthly

for the market. Roddy recognises this as a serious potential switch in long-term psychology towards commodities as it is a classic double-bottom price pattern whereby the reaction high has been exceeded and now, when it is being re-tested, it is acting as support for the market. There is clearly long-term demand for commodities underlying the market and the conclusion from a trading regime perspective would be that there is a high probability of commodities moving higher in an impulsive trending, higher volatility regime over the next few years. Roddy decides to factor this into his overall asset allocation decision process, which includes an analysis of the fundamental developments in commodities, and finally comes to the conclusion that an overweight position in commodities would seem to be the appropriate asset allocation decision to make in the middle of 2003. His asset allocation note to investors includes a reference to this psychological development in the market, and also highlights the risks to this decision being wrong, including a reference to the fact that if the CRB index started to move back down below the old high level it would eliminate the positive psychological analysis.

As things turned out, the decision to overweight commodities in July 2003 was a master stroke as the CRB index roared higher over the next few years.

THE DAY TRADER

Una is retired but concentrates full time on day trading to enhance her pension income. She concentrates her activity on just one market – the DAX (German equity market) future.

Figure 25.8 shows the chart that Una stares at all day when the DAX is trading. This particular chart shows the DAX Index future traded on the Eurex exchange expiring in June 2008 and using the 5-minute time fractal. The lower chart, showing the 20-period Bollinger Band Width, is her preferred measure of range expansion and contraction. Una's trading methodology is very simple. She only trades when she sees a breakout occurring from a clear contraction in price range (or drop in volatility). She does this in the expectation that there will be a move from a range-trading regime to a trending regime and, in this respect, her trading methodology is a classic volatility breakout method.

In the DAX future each tick of 0.50 is worth €25 and Una trades a maximum of 5 contracts. This means that she can take positions to the value of around €900,000, depending on the current index level, with only a fraction of that put up as margin with her broker; such is the power of leverage. Una's money management plan is to put her maximum position size on a volatility breakout, close out three of those if the break gives a 10-point move, move her stop-loss on the remaining two to her entry price, and then either get stopped on the other two or close them out at a 20-point move from entry. This is a classic money management method whereby the trader is taking some profit but leaving some potential for more profit at zero cost (because the stop has been moved).




Figure 25.8 DAX Index Future June 08 – 5 Minute Chart – 7 May 2008

At 2pm European time on 7 May 2008, Una notices that the DAX future has, once again, held the support zone of an ever-decreasing range that resembles a falling wedge. It could also be counted as a fourth wave consolidation for that day under Elliott Wave conditions. The Bollinger Band Width, not surprisingly, is extremely narrow, so Una is poised to take advantage of a range breakout into a trending regime. She does not, necessarily, know which way it will break but, because of the analysis, expects it will be to the upside.

In the event, the DAX future did break to the upside and it allowed Una to go long five contracts at a level of 7119. As per her money management plan, she sold three of those contracts at 7129 and the other two at 7139. This gave her a gross profit of €3,500.

THE OPTIONS TRADER

Alyson works for a charity during the day but in the evening she turns into an equity options trader. Alyson trades individual stocks that are traded on the OEX Index, which is the Standard & Poor's 100 Index of highly capitalised stocks for which options are listed. She decided to concentrate her efforts on these stocks only because they are liquid markets and 100 stocks give her more than enough opportunity and diversification. She gets home from work and spends an hour in the evening going through her charts of the stocks. When she sees an opportunity she wants to trade

on she puts her orders in with her broker the next day and, therefore, Alyson is not a short-term or day trader. Instead her investment timeframe is in the medium term of anything between two and eight weeks. She is flexible as to the strategies she is likely to put on but she considers her core strategies to sell directional volatility via call and put spreads and to buy volatility via straddles.

The date is 17 October 2007 and Alyson has noticed that Citigroup Inc. (C) is a potential trade (see Figure 25.9). Citigroup has been trading in a distinct range since July and, in fact, looks like a descending triangle with the lower boundary coming in around \$45. On this day Citigroup is trading just below \$45 and so a potential break of the triangle may be taking place. Also today there was a big surge in volume as the stock broke down through this support level, adding to the evidence that a sustained break might be occurring. The Bollinger Band Width on the daily time fractal is at a low level compared with history, and this is yet another clue that the odds on a trending regime in the next period are quite high.




Figure 25.9 Citigroup daily and Bollinger Band Width

Alyson decides that there is a good chance that Citigroup is about to enter into a trending period, and while she thinks the chances of that trend occurring to the downside are higher than the upside potential, she is not sure and decides that the best strategy would be to buy a straddle. She is analysing Citigroup on the daily time fractal, meaning that a trending move, if it were to occur, should happen

fairly soon but because she is wary of the time decay effect of owning options she decides to buy the straddles in the three-month tenor. She places an order with her broker to buy the \$42.50 January straddle for the next day at the open. The broker executes the order at the open and sees that the calls are priced at \$3.74 to buy and the puts are priced at \$2.00 to buy. Alyson buys 10 contracts of each costing her a total of \$5,740 (10 contracts of calls at $\$3.74 = \$3,740$ and 10 contracts of puts at $\$2.00 = \$2,000$). She also pays a total \$100 in commission to her online broker for the transactions. Her total outlay for the strategy is therefore \$5,840.

Over the next few days Citigroup continues to fall in price and the band width continues to rally from its low levels. Alyson sees that her trade is becoming profitable but decides to hold on to the position because she can't see from the price action any sign that the down move has exhausted. When Citigroup breaks \$40 on 1 November there is a big spike in volume, but the close on the day is towards the lower end of the range and so Alyson decides to hold. Bollinger Band Width is still pointing up and is not extremely high.

On 8 November there is another big spike in volume, but this time the price action has the look of a hammer bottom. The share price went down a lot during the day but closed the day up near the opening levels. The Bollinger Band Width is now quite extreme from a historical perspective on the daily time fractal, and at this point Alyson decides that she will close the trade.

At the opening on 9 November the trade is closed. The calls are sold for \$0.35 and the puts are sold for \$10.70. Again the commission amounts to a total of \$100. The loss on the calls is $\$3.74 - \$0.35 = \$3.39$ per contract and the gain on the puts is $\$10.70 - \$2.00 = \$8.70$ per contract. Therefore the gain on the options is $\$8.70 - \$3.39 = \$5,310$. Take away the \$200 paid in commission and the total net gain is \$5,110. She sold the options for a value of \$11,050 (calls for \$0.35 and puts for \$10.70) and taking away the \$100 commission the amount received into her account is \$10,950. Her initial total outlay was \$5,840, so she has made a net \$5,110 or 87.5% return on her initial investment.

A few weeks later on 7 December 2007 Alyson spots another opportunity in the options market, this time in HJ Heinz Co. (HNZ). Earlier in the year Heinz had reached a high of \$48 and then \$48.70, but that price zone had acted as a resistance area thereafter. Now, on the week ending 7 December the market had re-tested this price zone but three pieces of evidence came to light that made Alyson think the market may be near a trend exhaustion point. Firstly, the weekly candlestick was a spinning top and very close to being a doji. Secondly, the market had done five waves up on the daily time fractal, indicating that conditions were near right for a period of range trading. Lastly, the Bollinger Band Width in the daily time fractal was quite high by historical standards. All of these things made Alyson confident that topside volatility in Heinz was quite rich and that the market would either go sideways for some time or go down.




Figure 25.10 Heinz Co, daily and Bollinger Band Width

She therefore decided that a good strategy would be to put on a bear call spread where she sells a slightly out-of-the-money call and buys an even further out-of-the-money call for the same expiry date. On 10 December she sells 20 January \$50 calls for \$0.40 and buys 20 January \$55 calls for \$0.04. This results in a credit into her account of \$0.36 per contract or \$1,800. This premium that she collects is small compared to the maximum at maturity risk of the strategy, but Alyson feels that the trade is worth the risk because the trend exhaustion signals are so strong. We can see from Figure 25.10 that over the next few weeks Heinz did indeed fall back into a range-trading to down pattern and so by the maturity of the options in January both call options expired worthless and Alyson kept the \$1,800 she had collected as premium back in December.

THE TECHNICAL ANALYST

Rodney is a technical analyst for a commodity broker whose main responsibility is to advise the bank’s customers on the short-term technical developments within the commodity markets. It is March 2008 and Rodney is reviewing his advice over the past few months in the wheat futures contract. He has been advising his clients on

pattern and he was already advising his clients that a break of the upper limit around 695 5/8 would, in all probability, result in a swift, trending move higher.

On 21 November the range-trading regime broke and the price began a sharp advance over the next few weeks, reaching a high of 833 on 13 December. This range expansion had of course meant that the standard deviation band width had risen a lot and it had in fact reached the historical highs yet again. Rodney noticed that both the 12 and 13 December trading sessions were outside days, meaning that the highs and lows were outside the previous day and this, for technical analysts, represents a loss of control of the market by the incumbents (in this case of a previous uptrend, the bulls). This switch in psychology is a sign that the trading regime is changing, so Rodney had warned his clients to expect the advance to have been exhausted with the price either moving sideways for some time or down. As it turned out, the price of wheat moved sideways over the next few weeks resulting in a distinct range-trading regime that, by 11 January, could be labelled as a symmetrical triangular price pattern. With volatility being once again at the lows, as measured by the Bollinger Band Width, Rodney's analysis was that a break either way from this consolidation pattern should result in a strong trending move. We can see from the chart that the market broke higher above the resistance zone on 14 January and continued higher over the next few weeks. This recent activity in wheat, Rodney concluded, was a classic example of how rhythmically mean reverting the volatility cycle can be and how orthodox technical analysis can anticipate the expansion of volatility quite well.

THE LONG-ONLY EQUITY FUND MANAGER

Nicola is a long-only equity fund manager whose main responsibility is to manage a fund of UK stocks. Her fund is constrained by the fact that it has to be invested in stocks and can only ever have a maximum of 10% cash. It cannot short stocks. These rules force the fund to be compared against an index, in this case the FTSE 350 stock index, and so the fund can really only beat the index by anticipating sector rotation and relative value plays. Nicola has an investment process that looks at the fundamental analysis of various sectors but also takes into consideration the trading regime in a sector relative to the overall index in order to decide whether to be underweight, overweight or neutral of the sector in the fund. One of the tools she employs is the moving average envelopes which allows her to use an objective measure of whether the sector in question is in an uptrend, downtrend or neutral relative to the overall market, and because she is a long-term investor she uses the monthly time fractal to analyse this trading regime. It is May 2005 and Nicola is wondering how to weight the funds holdings in the media sector, a sector that has a relatively high ranking in weighting terms of the overall index and includes such companies as WPP Group PLC and Reed Elsevier PLC.

a July 2008 wheat futures contract since the middle of 2007 and, from a technical analyst's point of view, he is very pleased because this contract has exhibited classical trading regime behaviour that made his analysis valuable.

Figure 25.11 shows the bar chart of the July 2008 wheat futures contract using the daily time fractal and, below, the Bollinger Band Width. At the start of October 2007 Rodney noticed that the sharp rise in the price of wheat had resulted in the Bollinger Band Width reaching historically high levels, and by 9 October he had noticed that this volatility expansion seemed to be leading to volatility contraction with the price of wheat being contained within an upper level of $695 \frac{5}{8}$ and a lower level of $648 \frac{3}{16}$ (the wheat futures market as is common with many markets still trades in fractions rather than decimals). At this point Rodney's advice to his clients was that a period of volatility contraction was quite probable and, as time went on, this turned out to be the case with the market's psychology in such a state of uncertainty that a distinct rectangular price pattern emerged. By 20 November the price had spent so much time in the distinct trading range that the standard deviations (as measured by the band width) had collapsed down to historically low levels. By this stage Rodney's technical analysis had identified the contraction of volatility as a rectangular price pattern




Figure 25.11 Wheat (July 08) and Bollinger Band Width

indicating that the trading regime may be switching from a range-trading regime to a strong downtrending regime where the media sector will be underperforming the FTSE 350 Index as a whole. If this fits with Nicola's other assessments of the media sector, she should be encouraged to position her funds underweight the media sector. We can see from the chart that if she had chosen to do that, the fund would have benefited over the next 19 months as the Media Sector Index rose only 13% relative to the FTSE 350 Index rising by 29%.

THE MARKET MAKER

Paul is a foreign exchange market maker in the spot market. His main responsibility is to manage the GBPUSD or "cable" book and this has his attention throughout most of the day. Paul's business is very, very short term in nature. He does not have the mandate to keep risk on his book for very long and must go home flat at the end of the day (that is, he cannot go home with any open positions on his book). Paul knows there are times when the markets will be quiet and times when it will be busy. He knows there will be a lot of activity and "noise" around major economic data releases from the UK or the USA as these will affect people's perceptions of the outlook for the GBPUSD exchange rate. He also knows that in Tokyo hours the market will be generally quiet because there is usually not a great deal of market volume around at that time, especially in the GBPUSD exchange rate. He knows this through market experience and will factor this knowledge into his decision-making process when he goes about his market-making business. Paul's goal is to make a "turn" on his market-making activity by selling his inventory (the currency he buys when customers sell it to him) at a slightly higher price than where he bought it and, depending on market conditions, the market price will either give him a profit or loss quickly or slowly.

Market makers deal "inside the spread", meaning that the market price for them is a very tight difference between where they can buy and sell. The price they "make" to customers, however, is slightly wider than this and this concept is the same as any other business whereby different prices will be charged at different levels along the supply chain. The market maker's job is to provide liquidity to customers who want to buy and sell their currencies, and for assuming this relatively risky role he has to be compensated. For instance, Paul might quote a price for a fund manager customer in the GBPUSD exchange rate with a bid-offer spread of 3 pips, for example 1.9620 to 1.9623. This means that the fund manager can buy GBP at 1.9623 and can sell GBP at 1.9620. Let us assume that he decides to sell GBP to Paul at 1.9620. Paul is now long of GBP (and short of USD) at a cost price to him of 1.9620. However, because of Paul's status as a market maker he has access to prices that are inside the spread and he can therefore sell the GBP that he just bought from the fund manager at a price just higher than 1.9620. He might be able to sell his GBP to another market maker at a price of 1.9621 and in so doing he makes a profit of 1 pip on his trade.




Figure 25.12 UK Media sector relative and MA Envelopes

Figure 25.12 shows the line chart monthly closing prices of the UK FTSE 350 Media Sector Index relative to the parent index, the FTSE 350. It also shows the moving average envelopes using a 24-month moving average and 5% bands. Nicola uses the logic that if the price relative is within the envelopes then this constitutes a range-trading regime and therefore the sector's performance should be stable relative to the index. By the same token, if the price relative is above the upper envelope, the sector should be outperforming the index; and if the price relative is below the lower envelope, the sector should be underperforming the overall index.

Nicola has looked at the history of this trading regime indicator and, while there have been a few false starts, the indicator did a reasonably good job in identifying periods when there was persistence in the regime. Back in the early 1990s, for instance, the indicator identified that the media sector price relative was in a strong uptrending regime when the media sector was outperforming the FTSE 350 Index between February 1992 and September 1994. In that time the Media Sector Index went up by 47% compared with an increase in the FTSE 350 of only 22%. More recently Nicola remembers the period between March 2001 and June 2003 when the media sector underperformed the overall market with the media sector falling some 44% compared with the 350 falling only 27%. The moving average envelope method of identifying trading regime picked that period out as a period of strong downtrend for the price relative, and so Nicola was underweight the media sector during that time. Since June 2003 the price relative has stayed pretty much within the envelopes, meaning that Nicola has maintained quite a neutral sector weighting in the fund. There were a few times when the price relative moved outside the envelopes, but it moved back in very quickly thereafter, allowing Nicola to cut her losses promptly. In May 2005 the price relative has popped out of the bottom envelope, which is again

However, what if Paul factored in an analysis of trading regime into his decision of whether to “turn” his inventory now or later? If, for instance, the analysis suggested that a period of range trading was more likely, he might be inclined to turn his inventory quickly, but if the analysis suggested that a big trending move had a decent probability of occurring then, depending on his position, he might be inclined to hold his inventory and take advantage of the move in price. To analyse the trading regime that is applicable over his instantaneous and reactionary time horizon, Paul will want to look at the microscopic time fractals such as 1, 2 or 3 minute charts.




Figure 25.13 GBPUSD 1 minute – 6 May 2008 and Bollinger Band Width

Figure 25.13 shows the chart of GBPUSD in the 1-minute time fractal on 6 May 2008 with the associated Bollinger Band Width below. We can see that at 12.50pm London time the trading or volatility analysis would have encouraged Paul to make prices that skewed his trading book to accepting USD rather than GBP because, with the band width being so high, the chances were that the previous micro term uptrend was reaching an exhaustion point. So it proved, with the market then falling in a distinctive a–b–c correction before rallying up to new highs at 1.9745 and more trend exhaustion.

By 2.30pm Paul would have noticed that a distinctive rectangle had formed and, with the Bollinger Band Width being appropriately low, he would be expecting the market to enter a period of trend when the range was broken. So at 2.30pm he would

want to skew his pricing more to accepting GBP rather than USD as he saw the upside resistance start to break.

The use of this volatility analysis would have helped Paul to skim a few more pips out of his market making activity for that day and he could retire down to the local pub with a few more “war stories” to tell.*

THE STOCKBROKER

Cameron is a stockbroker who advises private clients on UK stock market investments. He earns his money from commissions that his clients pay when they deal through his firm and so, for Cameron, large volumes of transactions are good




Figure 25.14 Next PLC weekly

* In actual fact this case study is becoming rapidly out of date because the major foreign exchange rates now are being priced using algorithms rather than humans owing to the massive and predominant volume that is going through electronic trading systems. Human market makers in major FX still exist but really only act as back up to the computer. Nevertheless, an appreciation of volatility analysis like this can also be factored into a “pricing engine” to skew bids and offers one way or another. The only difference is that the computer doesn’t go down to the pub after work to exchange market-making “war stories” with other computers!

news. When there is volatility in the markets that usually equals large volumes, and Cameron looks for situations where there is likely to be volatility breakouts in UK stocks. The date is 15 December 2006, and Cameron has been studying the charts of various UK stocks when the share of Next PLC, the high street retailer, catches his eye (Figure 25.14).

He notices that the share price has rallied from a low of 1245 pence in October 2005 in five clear moves up to a high of 1747 pence in early January 2006. This five wave move was a classic Elliott Wave move and so the fact that the market had then entered a long period of sideways consolidation was not unexpected. In fact, the long period of consolidation had turned out to be a distinct rectangular pattern with the upper and lower bounds clearly defined albeit as zones rather than precise price levels. On the week of 15 September 2006 the share price of Next broke higher above the resistance zone of the rectangle and did so on a massive increase in volume. This was a sure sign that the market had broken the range-trading regime and, in all probability, would be likely to enter a trending regime. However, a doji candle week on 6 October suggested that a further push past the 1975 pence highs may be delayed and, as it turned out, the market moved sideways to lower over the next few weeks coming all the way back to the breakout zone around 1747 pence. It is this zone that now, viewed from the week of 15 December 2006, seems to be holding a support and from Cameron's point of view this is very useful and interesting information. A break higher above an area of range trading and then a re-test of that zone that holds as support, is a very big clue that there is strong underlying demand in the market and that the probabilities are skewed towards the market entering a strong uptrending regime in the period directly ahead.

Therefore Cameron decides to advise his clients to take a long position in Next PLC and those that acted upon this advice would have been rewarded with a move up in the share price in a strong trending regime over the next few months with the share price reaching a high of 2504 pence (43% up from the breakout zone) on the week of 30 March. It was on this week in fact that another development occurred that would have encouraged Cameron to, probably, advise that the up move had exhausted itself. The market opened at 2227 pence, reached the high of 2504 but then closed the week all the way back at 2250 pence. This is a classic shooting star in Japanese candlestick analysis and indicates after an uptrend that the underlying market psychology has changed from bullish to neutral. Being a trading regime analyst Cameron would have noticed this and advised his clients that the uptrend was probably exhausting and that a period of range trading to weakness was probably on the cards.

26

There is Still No Holy Grail

“After an abominable round of golf, a man is known to have slit his wrists with a razor blade, and, having bandaged them, to have stumbled into the locker room and enquired of his partner, ‘What time tomorrow?’ ”

Alistair Cooke

No matter how hard some people try to search for the holy grail of market trading or investment techniques that perform in each and every time period, they will never find it because none exists. In my opinion financial market analysis is probably best approached from the perspective of the social scientist rather than the perspective of the science of physics or engineering. The markets do adhere to a pattern but with that pattern being evolutionary, fractalised, chaotic and complex it is virtually impossible to model in the sense of creating a hard computer program that would take account of the entirety of this evolutionary process. As I have said, however, that does not mean that mechanical trading or investment programs do not have any value. They do. In fact, in a lot of respects mechanical programs are vastly superior to discretionary programs but mostly because they eliminate the emotional aspect of investing in the markets and allow decisions to be made in a very objective manner. This is the main benefit of using mechanical models to invest in the markets as far as I am concerned. The markets evolve and change over time but this does not mean that one successful method of trading the markets dies forever. Rather, this evolving change has a repetitive nature, which is manifested in the changing cycle of volatility. In this respect the markets evolve and change, but the change is cyclical rather than permanent.

Is trading regime analysis just a fancy way of saying that something that has had a big, fast move will probably slow down and either reverse or trade in a range, or something that has not moved for some time will probably have a big move soon? Well, to be perfectly honest, yes it is. However, in my opinion market participants constantly overlook this simple, yet powerful, observation and one reason

could be that market participants get sucked into the psychology of the current trading environment. We all know how difficult it is for us not be sucked into a move *after* it has already moved a lot, and we also know that it is easy to dismiss putting positions on or get disheartened being “whipsawed” in a quiet market. I think this is why people forget about the mean-reverting nature of volatility. If something is currently volatile somehow the current volatility is extended into the future with people expecting the move to continue forever, and this of course is all tied up in the psychology of the market cycle. The most bullish noises are at the top and the most bearish noises are at the bottom. Similarly, if a market is extremely quiet and range bound, we somehow all dismiss the potential the market may have in entering a large move. The market story is quiet, not much is happening and so we project that into the future.

TECHNOMIC ANALYSIS: A HYBRID APPROACH

As I said from the outset I am a trend follower. I follow trends. That is what I do. I feel comfortable following trends because I know it will guarantee me to be on the correct side of a big, sustained move in the market. I don't like it when markets chop around because that is when trend followers lose money. I know I will make money when markets trend and I know I will lose money when markets do not trend. These are the facts of my life.

As a trend follower, therefore, what I want to do is to maximise the winning streaks and minimise the losing streaks, and I can do this by adjusting position size based on a study of the trading regime or volatility conditions of the market. By increasing position size when volatility is low, I aim to take advantage of an increase in volatility that will manifest itself in strong trends, and by decreasing position size when volatility is high, I aim to guard against increased choppiness or range trading in the markets when trends will be weak.

Most trend followers are systematic in nature with regard to entry and exit rules, and beyond that most of us are systematic with regard to position sizing. The Turtle method, for instance, is fully automated with hard rules for how much risk to put and take off under certain market conditions, and most trend followers believe that a fully mechanical system is desirable, not least for the fact that only a fully automated system that includes position sizing will be able to be back-tested. Furthermore, only a fully back-testable system will be able to give the trader or investor the fullest confidence that this system will provide results over the long term.

However, surely it is worth asking the question, if only from a philosophical point of view, of whether it is possible to combine a systematic, mechanical entry and exit method with a discretionary evaluation of position size and risk. Is it possible to be a sort of discretionary trend follower whereby the direction of your investment is mechanised, but you have a discretionary control over the size of the position? This methodology is attractive to those of us who not only like the discipline and the

benefits that are provided by a guarantee to be on the big trends, but also want to be able to protect gains in times of trend exhaustion or limit losses in times of range trading, although we are not entirely convinced that an automated process of risk sizing is any better than a subjective one. This, of course, is anathema to dyed-in-the-wool systematic trend followers who would say that a discretionary position sizing methodology would only be prone to the natural human tendency to cut winners too early and run losers too long. Well, yes and no. I can see the point that any discretionary process is going to be subject to the human psychological condition of cognitive dissonance, but I can also see that an iron-clad discipline of investment direction would greatly minimise the extent of what could go wrong in that regard.

That is, if the investors are committed to following the trend, then they will have to be short when the market is going down and long when the market is going up. Thus, they will always be on the right side of a big trend in the market and never on the wrong side of one. The losers will be cut because there will have to be a position flip if the price dictates, and the winners will be run because there will always have to be some position on the trend. The only question remains, “how much risk will be on the trend?” and that, of course, is the real deciding factor in whether a method or process delivers results.

The trading regime analyses we have discussed in this book go a long way to helping the trend follower to decide on position size. Some of the analysis techniques are already factored into fully automated processes such as the Average True Range with regard to the Turtle method, and a million other trend followers will have some sort of automation based on any, all or none of these techniques for coming to a decision on size. Personally I prefer to use a discretionary analysis of volatility when deciding on position size because I feel it gives me the flexibility to utilise not only a multi-fractal but also crucially a multi-analyses approach. So rather than just look at the Average True Range using one time fractal I want to be able to examine the vast array of volatility measures and indicators over a few time fractals in order to gain as full a picture as possible of the potential trading regime that the next period will hold.

Moreover why would we, as trend followers, want to limit ourselves to a maximum position size? There may be times when the opportunity to really put our foot on the accelerator would be extremely beneficial, and if we know that we have the safety stop of our entry and exit system, then we should be more comfortable taking that risk. Of course it is a double-edged sword and if a trend system got chopped around with an outsized position, the draw-down would be nasty but, in my opinion, it might be worth taking that risk for the reward of extraordinary profits.

What I am about to say now will come as a shock to anyone who has read through this book from the start. At the beginning of the book I pointed out the differences between the fundamental and technical schools of thought with regard to market analysis, and went on to point out that fundamental practitioners tend to be more prone to the psychological flaws that we all have as human beings in the market. The so-called fundamentals are always most bullish at the top of a long trend, and always

most bearish at the bottom of a long trend, and that is when a practical understanding of technical market analysis saves the investor from getting sucked into a market that is about to reverse the trend. However, even I have to admit that there are times when the perception of the fundamentals line up with the early trend in price, and it is at these moments that the most powerful move in a market can be made. In Elliott Wave analysis this is identified as the third wave of a third wave, and this always tends to be the most impulsive. In the stylised market psychology cycle presented earlier, this would be the equivalent of when the current psychology is one of happiness (on the way up when the news is improving) or sadness (on the way down when the news is deteriorating). These are the times when the fundamentals and technicals are lined up together, so perhaps we, as trend followers, should be prepared to take account of this and greatly increase our position sizes during these times. Using this “TecHonomic” (technical and economic analysis) approach we could blend a mechanical discipline with discretionary flexibility to alter the size but not the direction of the risk.

Whoa. Take me out and whip me. Did I just say that? I know to technical analysts and mechanical trend followers this is sacrilegious but I am quite comfortable at some points to incorporate an awareness of the so-called fundamentals (and especially a reading of the socioeconomic conditions) into my judgement on position size. The reason why I feel comfortable doing this is because, if I am wrong in my assessment of the fundamentals (which I am perfectly happy to concede is a very subjective guess), then I know that if the trend turned against the position I would be stopped out and would be forced to re-assess. The entry and exit rules of the methodology force an iron-clad discipline which, in the final analysis, might not contribute to the best performance (that being more dependent on position size) but will mean that disasters will be avoided. And avoiding disasters is one of the absolute keys to long-term success in the markets. To use a final golfing analogy, increasing position size on the basis of discretionary analysis is the equivalent of using a driver off the tee on a tight par 4 because you think this will enable an easier shot to the green, but when you hit it into the rough, you chip out back on the fairway in order to have a chance at a par and certainly avoid a disastrous triple bogey. You do not compound the problems by trying one risky shot after another. You do not do a Jean Van de Velde!*

Therefore, by having the rule that, wherever the trend is pointing, the investment must be positioned that way no matter what my fundamental view, gives me the

* Jean Van de Velde is a French golfer who when leading the 1999 British Open at Carnoustie, Scotland, with one hole to play and needing only a double-bogey six to win, used a driver off the tee and hit into the rough. Rather than chip out and play conservatively from that point on, he decided to go for the green but this shot landed in heavy rough. His third shot landed in a water hazard, the Barry Burn, and after considering playing the ball from the water he eventually decided to take a penalty drop, but his now fifth shot landed in the greenside bunker. He played out of the bunker and holed a courageous six-foot putt for a triple-bogey seven, dropping him into a play-off for the title. He lost the play-off. His decision making on the 18th hole is a classic example of the short-term madness that extreme cognitive dissonance can provoke.

discipline required to never let my fundamental view of the world dictate the *direction* of the position. My fundamental view of the world only ever affects the *size* of my investment decisions. As an example, if I think that the US Federal Reserve are going to cut interest rates more than the market expects, and that the US economy will be underperforming the other major economies in the world, you may think that I should be positioned short of (or underweight) the US dollar. However, what if the current trend of the US dollar (as measured by the Supply Demand Index or other such trend-following methods) is up? Should I go short the US dollar now? What if demand for the dollar, which is currently outstripping supply of the dollar, continues to expand at the expense of supply? To go short the US dollar *now* would be to fight against the current supply and demand conditions in the market. Would it not be more prudent to wait until the trend of the US dollar has turned down to go short the US dollar? At that point I would be happy that my fundamental view of what I *think* will or should happen is aligned to the technical supply and demand conditions in the market, reflecting what is happening. Moreover, when the trend of the US dollar does turn down and is now aligned with my fundamental view, would it not be better to short that US dollar in a large size relative to my normal investments? After all, I should be confident that the market as a whole is also beginning to anticipate a dollar decline at this point, and so perhaps my fundamental guess will turn out to be right.

If, on the other hand, the trend of the US dollar continues higher, what is the most prudent thing to do? I don't *expect* the US dollar to continue to go higher so, if not short, should I not be neutral? I could be but what if the US dollar continued to trend higher and higher for a long time? If I was neutral I would not benefit from this rise and I might be left thinking that it was a missed opportunity. Would it not be more *prudent* to actually go *long* the US dollar, even though my fundamental view is *bearish*, and to reflect the fact that I don't "believe" in the dollar rising only go long in a small size? In this sense I could be said to have hedged myself against my fundamental view being wrong.

This hybrid approach to trend following, where we allow a fundamental bias to affect only the position size ingredient in the investment decision, can therefore give the process a reasonable amount of flexibility while at the same time having an iron-clad discipline behind it. Using a motor car analogy, the steering wheel of the car will always point in the direction of the preferred fractal trend to follow but the accelerator and brake pedals are controlled by elements such as higher or lower fractal trends, trading regime analysis and, yes, even fundamental analysis. This, I believe, is a pragmatic approach to trend following that allows the analyst or trader to express his discretionary views in the context of a hard-wired process.

I am a trend follower in my soul. I believe that "prediction" of the future is one of the silliest and most futile exercises that we can ever undertake as human beings. How can we possibly think we might know what will happen in the future? I am willing to accept that a prediction of the very near future is much easier relative to a prediction of the far future, but it just seems silly to me that we spend so much of our

time thinking about what will happen in the future when it should be obvious to any student of life in general, and of the markets in particular, that anything can happen at any time – and when the black swan, fat tail events happen they generally tend to have been unforeseen (at least by those not following the trend!). “Anticipation” of the future is an altogether different thing and, being a trend follower, I believe that to truly take full advantage of change it makes sense to follow trends (in the markets, in society, in life) even though we cannot always understand why these trends are occurring. Understanding is one thing, acceptance is another, and a trend follower accepts reality for what it is. All we have is the here and now. This moment. This price. No one, no matter who they are, knows what the next price will be. But if we are prepared to let ourselves give up control of worrying about what the next price will be and simply accept and *act* on it, then we are guaranteed to take advantage of change anytime it occurs. We may not understand at that precise moment why the change is happening but we will be guaranteed to take advantage of it and manage our risk accordingly.

CONCLUSIONS

This book has been about a pragmatic approach to investing or trading in the financial and commodity markets. Whatever your investment or trading bias is, having an understanding of the trading regime or volatility conditions of the market is vital to the success of your overall game plan because you then have an edge in deciding crucial aspects of the process, such as position sizing and money management. The irony, for someone with my style bias, is that the market’s cycle through the different states of regime or volatility is much more of mean-reverting phenomena, and my trend-following thinking has to adapt to this. Trend followers accept reality for what it is and have to accept the reality that following the trend will experience losing trades most of the time as the process of constant change gives them a number of false starts. However, by analysing the volatility conditions of the market they can skew the odds of anticipating those false starts in their favour. **By analysing the trading regime of the market in question, they can improve the process of anticipating change by gauging the probability of change itself, not just the direction of that change.**

The one constant in markets and life is the ebb and flow of volatility. Anticipating those cycles in volatility can make all the difference to your enjoyment!

Appendix 1

Time Fractals and the Supply/Demand Index

Fractals refer to the phenomenon observed in Elliott Wave and Chaos Theories that natural life evolves in a series of self-similar patterns at each degree of evolution. Time fractals take this phenomenon to a market context by observing that a market evolves in a series of self-similar patterns at each degree of trend or each degree of time. We can see this clearly in the classic stylised Elliott Wave chart that each wave is subdivided into ever smaller and smaller waves of self-similar patterns.

Whether we agree with the self-similar theory of fractals or not, the fact that a market can be analysed at each time fractal level means that time fractals are extremely relevant. In fact they are the most relevant pieces of information for the market analyst because if someone asks me the question “What do you think about the market outlook?” I have to know the degree of time or trend he is talking about. Does he mean the outlook for the next six months, the next month, the next week or the next six hours? If he wants the outlook for the next six months, I will analyse the weekly time fractal chart, but if he wants the outlook for the next week then the weekly chart is quite useless. I would need to analyse the daily, or more probably hourly, chart for that particular analysis. If we think about it in trend-following terms, when we define a trend by the slope of a moving average, then a market can be in an uptrend in the monthly time fractal, a downtrend in the weekly fractal, an uptrend in the daily fractal and a downtrend in the hourly fractal.

Figure A1.1 shows the Nikkei 225 stock market index in Japan over the monthly, weekly and daily time fractals respectively as at 16 October 2007. Included in the charts is a 21-period moving average of the close and, if we define the trend as the slope of the moving average, we can see that on the monthly fractal the trend is up, on the weekly fractal the trend is down and on the daily fractal the trend is up. On each degree of time, therefore, analysis of a market can and does show different conclusions. In this case let’s say you wanted to be a long-term investor and therefore




Figure A1.1 Nikkei 225 monthly, weekly and daily time fractals

wanted to follow the monthly uptrend. Would you be long the market because the monthly trend is up, would you be flat because the weekly trend is down, or would you wait until the weekly trend turned up to go long the market if the monthly trend stayed up? If you were a shorter term investor who wanted to follow the weekly trend, would you be short the market because the weekly trend is up, would you be flat because the daily trend is up, or would you wait until the daily trend turned down to go short the market, given that the weekly trend stayed down? These questions can confuse market participants all the time and is one reason why trading psychology and behaviour occur in the way they do. An investor might have entered a long position in the Nikkei but, when he saw periods when the daily trend turned down, he may have started to feel quite nervous about his position. In fact, even though he is a long-term investor he may have become so nervous while observing the price action at the daily time fractal that he chose to cut his long position only to see the price turn back up and follow the direction of the monthly time fractal, which was up. As we know from the earlier section on psychology, people can get short term

when they want to be long term, and can get long term when they want to be short term. It is usually the winning trades that become short term when they should be kept as long term, and it is the losing trades that get long term when they should be kept as short term. In other words, people tend to cut winning trades too early and run losing trades too long when what they should be doing is the exact opposite.

An appreciation of time fractal analysis is crucial to having a consistent plan in the markets. If we start jumping around time fractals at a whim, then all we do is confuse ourselves and our analysis, which will then obviously affect our trading or investing.

THE SUPPLY/DEMAND INDEX

The more experienced I have become in the markets the more I have come to realise that this game of ours *should* be one of the simplest to play. However, as I personally am all too aware, the art of extracting consistent above-benchmark returns from the markets is clearly the most difficult endeavour any human being could hope to enter into. Nonetheless, a large part of my market analysis has essentially boiled down to the two elements this book is about. Either the market is going to be in a range-trading environment or it is going to be in a trending environment, and so realising this has allowed me to dismiss many of the gimmicks dressed up as new theories or “magic” models that appear to make money in every time period (at least in back-testing!). No, either the market is trending or range trading and in my experience most trend-following techniques will find the trend to a greater or lesser extent, and most range-trading techniques will be able to extract value from the range to a greater or lesser extent. In the final analysis, everything in the markets boils down to being either a strategy that benefits from increasing volatility or a strategy that benefits from decreasing volatility. It’s a long-volatility or a short-volatility world.

Over the years I have therefore come to the conclusion that my personal style bias, in the long time fractals at least, is towards trend following because I am a natural buyer of volatility and, moreover, that my preferred method of measuring the trend is using the venerable moving average, or two moving averages to be precise. As I have mentioned, in my early years in the markets I back-tested and optimised systems and moving averages before I finally realised that optimising systems and parameters is, in my opinion, just plain crazy. Sure we can optimise parameters to find the best fit for a set of data points in a time series but the elephant in the room is always the question of how often to re-optimize. To me it is just data fitting. So I decided a number of years ago that establishing parameters that were consistent and sensible was the most important thing to do, at least in relation to trend following. There is no magic. Moving averages provide a discipline, not a magic black box. I decided that my preferred brand of trend following was the very straightforward continuous moving average cross system where the system is long when the shorter average is above the longer average and vice versa.

As far as I am concerned the basic and most important “driver” of market prices is the balance between the supply of that market and the demand for that market. As I have mentioned previously, the “reasons” why people want to supply the market or demand the market are quite frankly secondary to the actual act of doing it. Someone could sell a huge chunk of shares that he owns in a perfectly good company because he got out of the bed on the wrong side and was in a bad mood. Or he might need to raise some cash due to outflows from his funds. The reasons for people acting in any market are so numerous that worrying about them and trying to figure out “why” people are doing what they are doing is, in my view, very futile. What is important, however, is working out where the balance of supply and demand lies and where it might lie in the future. Fundamental analysts do this by coming to a conclusion about the future prospects of the market based on their forecast of future variables that may or may not be correlated with the actual buying and selling of the market and, crucially, making a big assumption that the majority of the market power will also agree with that outlook. Technical analysts do it by analysing the current balance between supply and demand in the market itself and anticipating how that might change, given the current psychology of the market. Technical analysts are trying to work out who has control of the market, the bulls (demanders) or the bears (suppliers) and there are a number of ways this can be done. However, my preferred method is to use moving averages of the closing price, as this method is a standard way to smooth out the noise in the market.

I chose a moving average combination of 5 for the short one and 40 for the long one. There was no particular science involved in this and, in fact, it was probably an anti-science thing because I was so fed up concentrating on optimisation issues. I thought, why not choose something that feels sensible and can be used over all time fractals, and so I wanted an average that was short enough to be sensitive (but not too sensitive) and long enough to keep me in big trends (but not so rigid that it missed cycles). I picked 5 and 40 and have used these numbers ever since. To reiterate: there is no magic, only discipline.

The philosophy I use is very simple. When the 5-period moving average is below the 40-period moving average, this is showing that supply is dominating demand (therefore there is downward pressure on the market) and when the 5-period moving average is above the 40-period moving average this is showing that demand is dominating supply (therefore there is upward pressure on the market). This is what the difference between a short moving average of the market price and a long moving average of the market price tells me. It gives me the information of when supply and demand dynamics in the market are changing. If the market has been in a long downtrend (supply dominating) then the 5 will be below the 40, and will have been for some time. However, if demand starts to show itself in the market it will show itself first in the 5-period moving average, and if the 5 moves up above the 40, then this information tells me that demand is now so strong that it is dominating supply. A great many people use moving averages but I don't think too many actually stop to think about the information they give us, and I believe the supply and demand

information in looking at a short moving average against a long moving average is beautiful in its simplicity. Not only can this simple method tell us when supply is dominating demand or vice versa, but it can also tell us when supply or demand are at extreme points. This concept is measured by the distance the moving averages are away from each other.

Figure A1.2 shows the Structural Supply Demand Index for the US Dollar Index from 1990 until 2008. Structural refers to the fact that I am using the monthly time fractal for this chart and the solid black line is the monthly closing price of the US Dollar Index, shown on the scale on the right. The scale on the left shows the difference between the 5- and 40-month moving averages on the US Dollar Index and the area graph shows the chart of these values. When the chart is below the zero line it is showing that the 5-month moving average is below the 40-month moving average and that supply is dominating demand. Similarly, when the chart is above the zero line, it is showing that the 5-month moving average is above the 40-month moving average and that demand is dominating supply.




Figure A1.2 US Dollar Index (DXY) – Structural Supply Demand Index

We can see that when supply is dominating demand (and the ratio is declining) the market is in a downtrend, and when demand is dominating supply (and the ratio is increasing) the market is in an uptrend. This is simply what we would expect, given the fact that we are using the first derivative of the price itself (i.e. the moving average) and the movement in the index, as we can see, lags the price movement slightly. However, my belief is that periods of supply dominating and demand dominating can persist for a long time and, when they do, the market in question will experience a strong trend. The absolute key here is how often this persistence of supply or

demand occurs and, in actual fact, most of the time the periods of supply dominating or demand dominating will be relatively short, resulting in a choppy price action. However, when the minority periods of supply or demand persisting come along, the market's trend is so strong that it generally compensates, and more, for this period of choppy price action, or "whipsaw" in the language of trend following.

The bottom line here is simply this: What I am presenting is nothing radically new; it is just two moving averages being used in a trend-following system to show when there is likely to be an uptrend or when there is likely to be a downtrend. However, by presenting it in this fashion I hope I can get the point over to the reader that what moving averages show us is the most important element in any analysis of any market: that is, the balance between supply of and demand for that particular market. Supply and demand, as I am sure the fundamentalists and the economists would agree, is the ultimate arbiter of whether a market will go up, down or stay the same. There is, in the final analysis, nothing else to worry about. This is why I use the Supply/Demand Index as my most important and basic analysis technique for gauging market direction. "The trend is your friend" is one of the all-time truths in market lore. Although, as I have demonstrated in this book, it is possible to identify trend exhaustion and take advantage of that; to continuously fight a persistent trend is one of the most soul (and account) destroying activities anyone could wish to undertake. The trend is your friend. Burn it into your psyche.

Using a simple and consistent technique like the Supply/Demand Index enables the analyst to easily gauge supply and demand dynamics across multiple time frames or fractals. This gives the analyst a "global view" of the supply and demand mix because it will show how different time fractals are interacting with one another to give ideas and signals to trade or invest.

For example, Figure A1.3 shows the Supply/Demand Index for the DXY (US Dollar Index) over the Structural (monthly), Cyclical (weekly) and Tactical (daily) time fractals. We can see that presenting the data in this way immediately gives the analyst a sense of where the market is on all fractals at any one point of time. This was snapped on 26 May 2008, and on that date it shows that supply of the US Dollar Index was dominating over all time fractals. Notice that the Tactical Supply/Demand Index had signalled that supply was starting to dominate demand on 19 February 2008. On that date, too, the Cyclical and the Structural Supply/Demand Indices were showing supply dominating demand. On this date in February, therefore, the analyst would have seen a situation where supply was dominating in the two higher time fractals and that supply had just begun to dominate demand in the Tactical time fractal. We have a market where downward pressure is prevalent in two larger time fractals and where it is just becoming prevalent in a lower time fractal. When something like this appears it is usually a very good sell signal for the medium-term trader because the probabilities are stacked in favour of the supply pressure. This proved to be the case here with the US Dollar Index falling from around 76 to around 71 a few weeks later.

Using Supply/Demand Indices in this way, therefore, can identify some quite profitable opportunities for the trader or investor. In fact, this method of looking




Figure A1.3 US Dollar Index (DXY) – Supply/Demand Index

at two or three time fractals in order to stack the odds in favour of a trade is one that is well known and probably best described by Alexander Elder in his book *Trading for a Living*. Elder referred to his method as “triple” or “double screen” and he used slightly different techniques, but the philosophy is the same. The priority is to identify where the larger scale pressure (or trend) is in relation to the market and this can be determined by looking at the longer time fractals. The analyst can then look at the shorter time fractals to identify when there might be an extension of that pressure, thus producing buy or sell signals.

The beauty of the method is that it can be very flexible and used by long-term, medium-term, short-term and even day traders. For instance, look at the Cyclical Supply/Demand Index in Figure A1.3. Around April 2006 the Cyclical Index started to show that supply was dominating demand in the US dollar and this fact in itself is interesting. However, if we look at the Structural Supply/Demand Index for around the same time we can see that, although the supply pressure had weakened considerably from its extreme point around 2004, the overall supply pressure had still not abated (the index was still below zero). So when the analyst noticed that the Cyclical Index was showing that supply was dominating then, allied to the fact that the Structural Index also had supply dominating, it was a very big clue that the downtrend in the US dollar that had started in 2002 was not yet over. If the trader or investor had used this method to enter into a short US dollar position at that time, by the start of 2008 and obviously depending on their money management strategy,

he could have been sitting on the correct side of a 20%+ move in the currency market! This, of course, is far, far easier said than done (a fact to which I can testify!) and we all know that the best trader or investor in the market is Harold Hindsight. Nevertheless, even although I do not use the interaction between the time fractals of the Supply/Demand Index as a mechanical trading system, preferring to use it as a discretionary tool for position sizing, the method seems to offer some value for picking high probability occasions of trend extension. It is not infallible, nothing is, but it makes logical sense to me.




Figure A1.4 S&P 500 Index – Supply/Demand Index

Figure A1.4 shows the S&P 500 equity index in the United States of America and the relevant Supply/Demand Indices. We can see that the primary Supply/Demand Index was showing demand dominating all the way through the mega bull market of the 1990s turning to supply dominating in the downswing and then back to demand dominating from late 2003. What is interesting from this long-term index is that the power of the net demand in the upswing from 2003 had not, at the time of writing at least, come anywhere near the power of the net demand during the bull market of the 1990s even though the S&P Index itself had traded back up to the levels seen at the peak back in 2000. This is a sign that the underlying strength of the market is not as powerful as it was in the last upswing and is therefore a warning sign for bulls.

In fact, the S&P started to fall from re-testing the 2000 peak in 2007 and a study of the Intermediate Supply/Demand Index shows that supply was dominating from around November 2007, having admittedly only very briefly dominated in August

2007. Any long-term investor who was long the stock market towards the end of 2007 and was using the methodology of fractal based supply/demand indices, would (should) have been reducing his longs at the very least. We can see from the Minor Supply/Demand Index that supply started to dominate in early 2008, well before the big falls in the market in the middle of January which caused a bit of a minor panic at the time.

Therefore, using a simple study of supply and demand in the market can turn out to be quite useful in enabling investors and traders to effectively manage their positions on a dynamic basis, thus being able to either cushion the blow of trend reversals or take advantage of movements in the direction of the bigger trend.

AN INVESTMENT PROCESS IS SIMPLY A LONGER TERM (SUBJECTIVE) TRADING SYSTEM

A lot of fundamentalist analysts deride those that follow a systematic approach to investing or trading on the basis that those systems tend to be biased towards using price data as the main input. Even for those systems that use fundamental data inputs for making buy and sell decisions, there is criticism that they are inflexible and therefore not as useful as a discretionary methodology. As I have alluded to elsewhere, while I agree that a purely systematic approach can be quite dangerous at times when the human psychology responsible for managing that model is flawed, or when the model does not take into account market depth or liquidity (think LTCM), a systematic *approach* to investing is, in fact, what we are all trying to achieve.

The big institutional pension and mutual fund managers are given money to manage by their customers (such as large corporations or municipalities) and these customers will only consider giving them money if the fund consultants are satisfied that the fund manager in question has fulfilled certain criteria. What the consultants are looking for in a fund manager can very often be summarised as the three P's: performance, people and process. Performance is an obvious one because the customer will want to give his cash to a fund manager with a good track record of performance. In actual fact this is quite a debatable point because many people believe that, as performance tends to be mean reverting over time, you should actually rotate your allocations to fund managers by firing the top performing managers and hiring the worst performing managers. The philosophy is that the top performers over (say) the last three years will tend to start to underperform, whereas the worst performers over the same historical period will now tend to start to outperform. This logic is the same as that behind the Trend Following Performance Indicator. Periods of underperformance will eventually lead to periods of outperformance, and vice versa.

The second P stands for people, because fund consultants want to be confident that the people employed to manage and administer funds are (a) competent and (b) likely to be there for some years. A company with a high degree of staff turnover is, at least in the consultant's view, symptomatic of a company with problems in

Q: *What are the drivers of the market or stock in question?*

A: Numerous and varied drivers or reasons exist that cause market participants to act in the market by buying or selling.

Q: *What is changing in the market?*

A: The price is constantly changing to reflect the collective discounting of the future by the market crowd as a whole.

Q: *What is in the price already?*

A: All that is known or expected by market participants is in the current market price.

Q: *What are the triggers for the investment decision?*

A: When the designated short moving average of the price crosses the designated long moving average of the price from below, the decision is to buy/go long, and when the opposite happens the decision is to sell/go short.

This shows that an investment process is simply a set of rules that provide a framework within which the investment manager, trader or analyst can work. If it is consistency of investment process that is desired, then a process involving at least some form of technical market analysis (e.g. for the trigger decision) has to be better than using a purely subjective analysis.

SUMMARY

In this appendix we have looked at the application of fractal-based analysis on a market and how it can help to give the analyst a much broader understanding of the overall market conditions. Combining time fractals is, in my opinion, critical to understanding the context of the entire market, be that in terms of price or volatility direction. We have also examined how we can incorporate basic supply and demand time fractal analysis into an investment process in order to gain a higher degree of consistency.

the areas of staff morale and retention. As any business, in particular a service-based business, is nothing without its people, the consultants have to check this box as part of their work.

The third P stands for process, and this means the investment process. This is the area where fund managers can really differentiate themselves because it is the area in which their investment style or bias becomes evident. Some managers will have a bias to a fundamentally based process, some to a technically based process and some will have a process that includes both fundamental and technical market analysis. The most important aspects that a consultant will look for, however, is not whether the process is considered to be “good” in their opinion but that it is logical and, most importantly, *consistent*. Logic and consistency are therefore the watchwords for fund managers when they are thinking about their investment process, and the process will be written out and followed by the asset class teams involved in managing the funds. A very simple example of a fundamentally driven investment process might involve a list of criteria that have to be thought of or checked before an investment decision is made. That is, each asset class within the fund management operation might have to answer questions such as: What are the drivers of the market or stock in question? What is changing in the (fundamentals) of the market? What is in the price already? What are the triggers for the investment decision?

These questions provide a logical and consistent framework for an investment process, albeit one that gives a lot of scope for interpretation and meaning. The answers to these questions from a fundamental perspective will be quite subjective by nature, and one person’s answers might be different from another person’s answer. In this sense the investment process becomes quite fuzzy. For example, the so-called drivers of the market in question will surely be numerous and their relative importance will differ from one person to the next. I might think that the US dollar is being “driven” by the price of oil at the moment but my colleague might think that it is interest rate spreads. We can both produce evidence to support our case in the form of correlation coefficients. However, another colleague might say that this is just coincidence and that the real driver of the market is the central banks’ diversification of their reserves. Also, the answer to the question of what is in the price already is inherently subjective when viewed from a fundamentalist’s perspective. In some markets it is easy to identify what is “in the price”; for example, in the futures market of the Fed Funds 30-Day interest rate, there is an expiry date and a price, so the analyst can tell instantly the “price” of the Fed Funds rate at any given time. However, in most other ways the answer to the question of “what is in the price?” becomes a subjective debate which is impossible to resolve in any meaningful way. A stock might have had a big fall and I may think that the bad news that accompanied it is “in the price” already. My colleague, however, might disagree with that assessment and the problem is that there is really no way to resolve it.

Now consider a more technically minded investment process (for example, a trend-following methodology) that follows the same logic as above. The answers to the same questions now become very objective.

Appendix 2

Why Do Trend Lines Work?

Trend lines traditionally connect a series of highs or a series of lows on a time series price chart. Most people who come across technical analysis for the first time are immediately drawn [*sic*] to trend lines because of their simplicity and subjectivity. It is this simplicity in fact which makes trend lines some of the most valuable and important tools in a technical analyst's box. However, not many people stop to think about and communicate the logic behind trend lines, and this encourages the critics of technical analysis who would liken trend lines to "mumbo jumbo", alchemy or self-fulfilling prophecy!

Trend lines, in simple terms, show where supply and demand *probably* (not definitely) exist. In my opinion, therefore, they should be called *supply and demand lines* because this more accurately reflects the logic behind them. Supply and demand of the instrument in question drives the market price. Nothing else drives it. Not economic, technical or quantitative analysis. Not even economic or corporate statistics, as virtually everyone in the modern markets is conditioned to thinking, and what we spend most of our time talking about. It is, in fact, human beings that decide on the clearing price for the market (the market price) by interacting and haggling with each other. *It is how human beings act and react to information or analysis, not the information or analysis itself, that is the most important thing.* How many times has the market moved in the opposite direction after economic or corporate news to what conventional logic would have us believe?

Surely it is the job of the market analyst to identify where supply and demand of the *instrument itself actually* lies. Yet market analysts spend a lot of time analysing and espousing about where *they think* supply and demand *should* lie based on something that may or may not be related to the instrument (e.g. the dollar *should* fall, based on the current account deficit). This is where trend lines and technical analysis in general have a purpose. Technical analysis aims to discover how bullish or bearish the *market is as a whole*, not how bullish or bearish the *analyst* is. It can therefore be said to be unbiased.

Human beings have moods that swing. Our psychology is affected and shaped by events that happen to us. Is there anyone who can refute this point? If this is right and human beings decide what the clearing price of the market will be, it follows that the clearing price of the market will swing around with our emotions. It is because of this that the stock market is essentially a barometer of mass social psychology in a society. When markets move around, the moods of the participants in those markets are directly affected by the moving around of the current price. Anyone who refutes this has obviously not traded or invested anything in their life! We feel fearful and unhappy when the price is going against our position, and happiness and greedy when it is going with our position. Therefore, when the price moves, the actions of human beings will be dictated by the position they have and which way the price is going.

Take an example of an up move in price. If the market starts a move higher, those people who are long will feel happy but those short will feel unhappy. The market will move up until there are no more buyers and/or when the feeling of booking their profits overwhelms some of the longs. The market will then start to move lower and retrace some of the up move. How much retracement takes place depends on how much emotion is around by the market players. The people who are short and have just witnessed a move up will be relieved to see the price coming back down. Their emotions will be driving them to “get out” of the short position and they will buy back their shorts. Also, those that were flat when the market started to move up will be keen not to miss out if a bigger up move takes place and will be looking to establish long positions on this pull back in the price. In this way the market fall has created demand that has been driven by human emotions. If this process repeats a couple of times, the market has an uptrend. In fact, this emotional process starts to feed on itself. This is the real “self-fulfilling prophecy” in the markets but it is being driven by natural human psychology whichever brand of market analysis the market players are following, and not by one group of technical analysts! As a bigger move up happens, the people who are short are now even more fearful while the people who do not have a long position are even keener to establish one. So, again, when the price moves back down a little as buyers have run out and/or longs are booking their profits, the retracement in price is met with demand from people feeling relief (shorts) and greed (fresh longs).

As the market repeats this psychological process a series of higher highs and higher lows occur. All a trend line is doing by connecting the series of higher lows is identifying the strength (or the pace) of the psychological process. The strength and direction of the psychological process will dynamically change of course, and this is why trend lines will be broken and invalidated more often than not. However, a simple trend line is a very powerful and essential tool in identifying the direction and strength of the emotions of the market players.

Critics often attack the clear subjectivity in drawing trend lines. In my opinion, the analysts should not get too bothered about how precise trend lines should be. The markets are driven by psychology and psychology is, like economics, a social

science, not a hard science! The analysts should think of *zones* rather than precise levels when identifying potential support and resistance areas via trend lines. A good analogy is to think of drawing trend lines with a thick pencil. This allows for what some call “internal trend lines” where, in the example of an uptrend, the price might dip just below a trend line drawn with a sharp pencil but a thicker pencil would still have captured the low price! These, what at first glance look like “failed” trend line breaks when drawn with a thin pencil, often provide the strongest of signals because of the fact that the market has possibly taken out some standing orders (like stop-losses), but then realised that there is no more buying (in the case of a downtrend break) or selling (in the case of an uptrend break) to be done by anyone. By the economic law of supply and demand, the market has then to reverse and go back into the trend. Remember, all we are doing is identifying that a retracement in the trend has resulted in an emotional process by the market participants that manifests itself in (for an uptrend) a series of higher highs and higher lows. The best way to think of internal trend lines is as drawing a buffer zone around a trend line that takes account of general zones where supply or demand probably exist.

Lastly, think about this. If trend lines and technical analysis is a self-fulfilling prophecy, as some claim, how is it that trends can last for years when, if all the technical analysts in the world acted at the same time, that particular supply or demand would be absorbed into the market price very quickly?

Trend lines (or “supply and demand lines”) are not magic, “mumbo jumbo” or alchemy. The analysis is rooted in human psychology and crowd dynamics. It is the behavioural characteristics of human psychology that determine the clearing price of the market and trend lines are useful in helping to identify areas where these behavioural characteristics might occur.

Appendix 3

Examples of Trend Lines

HORIZONTAL SUPPORT

Let us examine the situation in Figure A3.1:

A: The sellers here during a downtrend expect prices to continue falling. (**SUPPLY**)

B: The market bounces a little and some of the sellers get scared. They hold on until new lows are seen and until they have a profit or flat. They then buy back their shorts here. Psychology is fear and relief. Also, new buyers sense value around the previous low. Psychology is greed. (**DEMAND**)

C: The market makes a bigger bounce and those still holding short positions are now very scared. The market comes back towards the previous levels where they had




Figure A3.1 EURUSD daily

sold and they take that opportunity to buy back their shorts. Psychology is fear and relief. Again, new buyers also sense value around the previous low. Psychology is greed. **(DEMAND)**

HORIZONTAL RESISTANCE

Let us now examine the situation in Figure A3.2:

- A: The buyers here during an uptrend expect prices to continue rising. **(DEMAND)**
- B: The market falls and buyers who got long near the previous high start to feel fear.
- C: The market rallies back to the previous high and those holding long positions who were feeling nervous and fearful see an opportunity to sell their long positions around breakeven. New sellers also come in and open shorts at what they perceive to be value (i.e. a previous high point). **(SUPPLY)**




Figure A3.2 EURUSD daily

UPTREND SUPPORT

Let us examine the situation in Figure A3.3:

- A: The market makes a bottom.
- B: The market rallies and then falls but support is found above the previous low. Demand has come in and outweighs any supply. There is short covering from




Figure A3.3 EURUSD daily

previous sellers who got short towards the end of the recent fall and now see an opportunity to cover. Their psychology is relief. Also, fresh buyers here are afraid they will miss out on a bigger rally. They don't want to wait any longer before buying. **(DEMAND)**

- C: The market rallies, falls back and finds support. The previous process repeats. The psychology of the fresh buyers who are creating the demand on retracements in the price is *fear* that they will miss out on a rally and *greed* as they see the market going up and want to get on board. So the price does not retrace all of the previous up move, a series of higher lows occur and a line can, sometimes, connect these lows. **(DEMAND)**

UPTREND RESISTANCE

Let us now look at the situation in Figure A3.4:

- A: The market makes a high above the previous significant high of early 2004.
- B: The market falls, then rallies again. The people who are still long notice that a high was made in early 2005 and, as the market moves above the old high, they sell. They sell their longs a little above the previous high because they fear that their profits will turn into losses. They cut winning trades. **(SUPPLY)**
- C: The market has continued to rally but pauses in early 2006 from the line before falling, in mid-2006, from the line. Sellers here have the same feelings as previously. They are holders of long positions who do not want to see their profits




Figure A3.4 S&P 500 weekly

turn into a loss. They cut winners. The market rises a bit past previous highs but then the fear of having a winner turn into a loser overwhelms the holders of long positions, so they sell their longs. (**SUPPLY**). The series of marginal new highs created by this psychology can, sometimes, be connected with a line.

DOWNTREND RESISTANCE

We can look at the situation in Figure A3.5:

- A: The market makes a top.
- B: The market falls and then rallies but resistance is found below the previous highs. Supply has come in and outweighs any demand. There is long closing from previous buyers who got long towards the end of the recent rise and now see an opportunity to close. Their psychology is relief. Also, fresh shorters here are feeling fear that they will miss out on a bigger fall in price. They don't want to wait until previous highs are beaten before selling and establishing short positions. (**SUPPLY**)
- C and D: The market falls, rallies back and finds resistance. The previous process repeats. The psychology of fresh shorters who are creating the supply on retracements in the price is fear that they will miss out on a fall in price and greed as they see the market going down and want to get on board. The price does not retrace all of the previous down move, but a series of lower highs occur and a line can, sometimes, connect these highs. (**SUPPLY**)




Figure A3.5 USDCAD weekly

DOWNTREND SUPPORT

Let us now look at the situation in Figure A3.6:

A: The market makes a low below the previous significant low.

B: The market rallies then falls again. The people who are still short notice that a significant low was made previously and, as the market moves below the old low,




Figure A3.6 USDCAD weekly

- they buy. They buy back their shorts a little below the previous low because they fear that their profits will turn into losses. They cut winning trades. **(DEMAND)**
- C: The market has continued to fall but rallies again from the line. Buyers here are having the same feelings as previously. They are holders of short positions who do not want to see their profits turn into a loss. They cut winners. The market falls a bit below previous lows but then the fear of having a winner turn into a loser overwhelms the holders of short positions, so they buy back their shorts. **(DEMAND)** The series of marginal new lows created by this psychology can, sometimes, be connected with a line.

References

- Borsellino, L.J. with Commins, P. (1999) *The Day Trader: From the Pit to the PC*. John Wiley & Sons, Inc., New York.
- Drobny, S. (2006) *Inside the House of Money: Top Hedge Fund Traders on Profiting in the Global Markets*. John Wiley & Sons, Inc., New York.
- Edwards, R.D. and Magee, J. (2001) *Technical Analysis of Stock Trends* (8th edition). St Lucie Press, New York.
- Frost, A.J. and Prechter, R.R. (1999) *Elliott Wave Principle: Key to Market Behavior*. John Wiley & Sons, Ltd., Chichester.
- Pring, M.J. (2002) *Technical Analysis Explained: The Successful Investor's Guide to Spotting Investment Trends and Turning Points*. McGraw-Hill, New York.
- Slater, R. (1996) *Soros, The Unauthorized Biography*. McGraw-Hill, New York.
- Smitten, R. (2001) *Jesse Livermore, World's Greatest Stock Trader*. John Wiley & Sons, Inc., New York.
- Steiner, R. (2007) *Mastering Financial Calculations* (2nd edition). FT Prentice Hall, UK.
- Taylor, F. (2003) *Mastering Foreign Exchange and Currency Options: A Practical Guide to the New Marketplace*. FT Pitman Publishing, UK.

Index

- Abbott Labs 353
- Absolute volatility analysis 57–8
- Actual market actions 22–3
- ADX 193–7, 213, 278, 291, 292, 302
- Alpha seekers 239
- Anticipation of price movements 24
- Anti-martingale method of money management 284–5
- Ascending triangles 81–6
- Asset allocator case study 373–6
- Assumed risk/reward ratio 284
- AUDCHF daily
 - and Bollinger Band Width 305
- AUDCHF hourly 305–7
- AUDJPY daily – rising wedge 121–2
- AUDUSD weekly 168–9
 - classic double bottom 114–16
 - rate of change 210–11
- Australian dollar exchange rate
 - vs Japanese yen, *see* AUDJPY
 - daily – rising wedge
 - vs US dollar, *see* AUDUSD weekly
- Auto Regressive Conditional Heteroskedasticity (ARCH) 54
- Average Directional Movement Index, *see* ADX
- Average True Range (ATR) 102, 193, 194, 231–5, 302
- Bank of England 33
- Bankers Trust 11
- Bar charts 79, 131
- Barings 33
- Bayesian probability logic 300
- Bearish engulfing pattern 136, 137, 138–9
- Bear Stearns 118
- Beauty contest concept 36
- Behavioural analysis 23–8
- Behavioural equilibrium exchange rates (BEER) 337
- Behavioural Finance 5
- Bet sizing 230
- Big Mac Index 337–8
- Black–Scholes Model 56
 - option-pricing model 263
- Blackstone 32
- Black swan events 98, 100
- Blaming 31
- Bollinger Bands 185–9, 236, 292
- Bollinger Band Width 70, 185–92, 205, 231, 236, 278, 291, 303, 307
 - cyclical fractal 300–2
- Bonnie Langford 263, 265
- Borsellino, Lewis 255–6
- Bretton Woods system 300
- Broadening patterns 101–3
- Brown, Gordon 33
- Bubbles 47
- Bullish engulfing pattern 136–8
- Buying power 41

- Call option 323, 324
- Canadian Dollar Index, *see* USDCAD weekly
- Candlestick 300
- Candlestick reversals 302
- Capital Asset Pricing Model (CAPM) 5, 53, 242
- Caremark Corporation weekly – falling wedge 124–6
- Carry strategies 14, 320–1
- Chaos Theory 5, 161, 178, 247
- Charting 199–206
- Chicago Board Options Exchange SPX Volatility Index (VIX) 44, 54
- Chicago Board of Trade 68
- Chinese Yuan 14
- Choice
 - of market to trade 310–11
 - of trading technique 315–18
- Cisco Systems weekly – falling wedge 126
- Citigroup weekly with volume – double top 116–17
- Close-to-the-money call option 325
- Clustering 54, 55
- Cognitive dissonance 25, 28–33, 41, 110
- Commodities Research Bureau (CRB) Index 294–5
- Commodity fundamentals 342
- Commodity Trading Advisor (CTA) 331
- Conditional independence 240–1
- Conoco Phillips weekly
 - and Bollinger Band Width 316–17
 - daily and slow stochastics 317–18
- Constant P&L 314
- Continuous systems 333
- Copper
 - monthly and Bolinger Band Width 190–1
 - Point and Figure 10 x 3 200–2
- Corrective waves 173
- Currency carry trade 320
- Currency fundamentals 342
- Currency overlay manager case study 367–72
- Cycles 5
- Cyclical fractal Bollinger Band Width 300–2
- Data fitting 282
- Data mining 282
- Day trader case study 376–7
- Delta 263–4, 271, 272, 273
- Delta neutral 271, 272
- Delta on the option 323, 324
- Demand curve 45–7
- Denial 31, 33
- De Niro 265
- Dennis, Richard 225
- Descending triangles 86–8
- Directional Movement (DM) 194–5
- Directional Movement Indicator (DM) 193–4, 195
- Direction of the market 18
- Dirty float 14
- Dirty Harry 265
- Discretionary market trading 14
- Dispersion of market 58
- Divergence 212–17
- Doji 131–6
- Donchian channels 225–37
- Dot com bubble 27, 31, 110
- Double tops and bottoms 113–18
- Dow Jones Industrials 44, 47, 155–9, 345
 - monthly and standard deviation 64–5
 - yearly – logarithmic scale 157–8
- Downtrend resistance 414
- Downtrend support 415–16
- Dow Theory 5, 80, 172, 286
- Duration overlay manager case study 372–3
- Dye, Tony 27
- East India Company 66
- Eckhardt, William 225
- Economists and fundamentalists 359–63
- Efficient market hypothesis 6
- Ego 34
- Elliott cycle 305
- Elliott Wave International 157, 162, 163
- Elliott Wave Principle 53, 75, 161–78

- Elliott Wave Theory 5, 63, 169–71, 300, 303
 - advantages and disadvantages 171–3
 - consolidation pattern 134
 - corrective waves 173
 - impulsive waves 173
 - overlap 173–5
 - rising wedge 124
 - theoretical basis 175–7
 - third of a third and third waves 169–71
 - triangle in 80, 88, 89, 90
 - vs Chaos Theory 247
- Engulfing pattern 136–9
- Enron weekly 98–101
- Equilibrium price 45–8
- EURGBP
 - daily 167–8
 - monthly – classic ascending triangle 81–5
 - weekly and 25-delta risk reversal 268–70
- Eurodollar futures contract weekly and slow stochastics 216–17
- Eurodollars 215
- Euro to the US dollar exchange rate, *see* EURUSD weekly
- EURUSD daily
 - moving average envelopes 180–1
 - time fractal 304–5
- EURUSD monthly
 - and Bollinger Band Width 313
 - trend following 315
- EURUSD weekly
 - Bollinger Bands 186–7
 - Bollinger Bands and Band Width 187–8
 - and Bollinger Band Width 303–4
 - classic symmetrical triangle 88–9
 - historical volatility vs standard deviation 60–2
 - implied volatility 56
 - MATE weekly 277–8
 - moving average envelopes 181–2
 - Nerve Index weekly 257–9
 - rate of change 208–9
 - time fractal 248–9, 314
 - trend exhaustion and 323
- Trend Following Performance Indicator (TFPI) 287–8
- and Volatility Smile Index 267
- vol of vol 65–6
- Exchange Rate Mechanism 312
- Expiry date of an option 253
- Extend U-shaped recovery 111
- Fair value 340
- Falling wedge 124–7
- Federal Reserve 215
- Fibonacci calculations 164
- Fibonacci numbers 283
- Fibonacci Sequence 175–6
- Flags 95–101
 - pennants 97–101
- Flexible game plan 165–9
- Forecasting 23, 24–5
- Foreign exchange markets 6–7
- Four-week rule 227
- Fundamental Equilibrium Exchange Rates (FEER) 337
- Fundamental market analysis 3, 4, 39
- Fundamental models 341–2
- Fundamentals 334
- Fund of funds products 318
- Game Theory 36
- Gamma 270–3
 - short 324
- Gann Analysis 5
- Gap theory 303
- Gazza 265
- GBPCHF Point and Figure 0.005 x 3 202–4
- GBPUSD
 - Purchasing Power Parity 339
 - weekly and relative value volatility 57–8
- Generalised Auto Regressive Conditional Heteroskedasticity (GARCH) 54
- Genetic algorithmic models 239
- Globalisation 66
- Gold
 - standard 90–1
 - as symmetrical triangle 90
 - weekly – classic symmetrical triangle 90–1

- Golden Ratio 177
 Goldman Sachs 312
 Google 47
 weekly – Bollinger Bands and Band Width 188–9
 Gould, Jay 120
 Great Crash 36
 Greater fool theory 38
 Greenspan, Alan 91
 Gross domestic product (GDP) 22

 Hammer bottom 129, 139–40
 shooting star 140
 Hanging man 140–1
 Head and shoulders 103–13
 reverse 109–13
 Historical volatility 50–3, 336
 Honeywell International Inc
 weekly candle 138–9
 Horizontal resistance 412–13
 Horizontal support 411–12
 House price inflation 348
 Hume, David 66
 Hungarian Forint 14
 Hurst Exponent 247–9

 Implied volatility 55–7
 Implied volatility curves 253–9
 Impulsive waves 173
 Indian Nifty daily – rising wedge 120–1
 Information Ratio 309
 Initial public offerings (IPOs) 31–2, 120
 Interest rate/bond fundamentals 342
 Internal trend lines 409
 Investor vs trader 309–10

 Japanese candlesticks 5, 129–41
 Japanese candlestick uncertainty 139–41
 JC Penney 350
 Johnson & Johnson 353
 JP Morgan 120, 200

 Keene, James 120
 Keltner Channels 183–4
 Keynesian liquidity trap 348

 Keynes, John Maynard 34–7, 54, 123
 Koh's Corporation 350

 Leading Indicator – Dow Jones Industrials 346–7
 Leeson, Nick 26, 33
 Limit moves 232
 Livermore, Jesse 199–200
 Long-only equity fund manager case study 382–4
 Long strangle
 case for 329–31
 strategy 326
 Long Term Capital Management (LTCM) 26, 33, 85
 Long volatility strategy 319

 Macy's Inc 350
 Managed float 14
 Mandelbrot, Benoit 54
 Market maker case study 384–6
 Market price 39–40
 Market psychology 84
 Market psychology cycle 74–6
 Marking to model 21
 Mark to market 21
 Markov chains 239–47
 Martingale system of money management 284–5
 Mask, the 265
 McDonald's Corporation (MCD)
 weekly – old high range 154–5
 Merck & Co. (MRK) 353
 monthly and volume 149–50
 Mexican peso 126
 Microsoft 101–3
 Momentum divergence 302
 Momentum indicator 213
 Moving Average Convergence/Divergence indicator 275
 Moving average envelopes (MAEs)
 179–84, 226, 276
 variations 183–4
 Moving averages 179–80, 275–6, 333
 MSFT 101–3
 My MATE (Moving Average Trading Environment) 275–9

- Nasdaq Composite Index – moving average envelopes 183
 Nash equilibrium 36
 National Stock Exchange of India 120
 Neural network models 239
 Nifty daily – rising wedge 120–1
 Nikkei 225 index 395–6
 NOKSEK weekly – old high range 152–4
 Non-continuous systems 333
 Non Farm Payroll number 69
 Northern Rock 33, 101
 daily – classic pennant 97–8
 Norwegian Krone Index (NOK)
 and ATR weekly 233–4
 see also NOKSEK weekly – old high range
 NZDJPY daily
 classic flag 95–7
 classic rectangle 92–3, 334–5
- OEX Index
 and ATR weekly 234–5
 monthly with volume 144–5
 monthly and Williams %R 220–2
 weekly and Williams %R 222–3
 Optimisation of parameters 282–3
 Option skews 263
 Options trader case study 377–80
 Option strange selling 325–9
 Option strategies 321–31
 trend exhaustion 321–4
 Organisation for Economic Co-operation and Development (OECD) 345
 Organisation of the Petroleum Exporting Countries (OPEC) 215
 Orthodox pattern recognition 79–127, 172
 Out-of-the-money call option 325
- P&L 314
 Pattern recognition 5
 Pennants 97–101
 Personal will 41
 Pfizer Inc 353
 Phi 177
 Phillips & Drew 27
 Plaza Accord 109
- Point and Figure (P&F) technique 5, 79, 199–206
 Position sizing 230, 231, 311–15
 Price discovery 20–1
 Price maker 20
 Price pattern market analysis 177
 Price taker 20
 Price zone 84
 Private swing trader case study 365–7
 Procter & Gamble weekly – classic ascending triangle 85–6
 Profit maximisers 6, 7
 Psychology
 cycle 74–6
 of range-trading regime 71–3
 of trending regime 73–4
 Purchasing Power Parity (PPP) 337, 338, 339
 Put/call parity 262
- Quantitative analysis 5
 Quantum Fund 312
- Range, definition 15
 Range trading 15, 149, 270–3
 Range-trading models 239
 Range-trading regime, psychology of 71–3
 Rate of change 360
 Rate of change indicator (ROC) 207–12
 Rationality 40–1
 RBS Regime Switching Model based on Markov Chains 246–7
 Reaction low 114
 Real Effective Exchange Rates (REER) 337
 Rectangles 91–5
 Relative interest rate spreads 356–7
 Relative Strength Index (RSI) 193
 Relative value volatility analysis 57–8
 Relative volatility analysis 57–8
 Resistance zone 84
 Reverse head and shoulders 109–13
 Rising wedge 119–24
 Risk budgeting 318–19
 Risk reversal 262, 269

- Risk-to-reward ratio 325
- Risk-taking 11–12
- Rothschild, Lord 37–8, 106
- Round numbers 155–9

- S&P 500 Index 132, 133, 146, 147, 148, 349
 - and Bollinger Band Width 349–50
 - CNX Nifty 120–1
 - doji trading range 132–4
 - monthly – rate of change 209–10
 - reverse head and shoulders 109–11, 112–13
 - supply/demand index 402–3
 - US Homebuilders sector 349, 350
 - weekly and volume 146–9
- Sector rotation and relative value 343–57
 - relativity and trending behaviour 343–56
- Self-fulfilling prophecy 343–4, 345, 408
- Selling power 41
- September 11, 2001 147–8
- SGDJPY
 - and Bollinger Band Width 307–8
 - hourly 308
- Sharks 83
- Sharpe Ratio 309
- Short gamma 324
- Short strangle strategy 326–7
- Short theta 325
- Short volatility strategy 319
- Singapore dollars priced in Japanese yen, *see* SGDJPY
- Skews 266, 268–70
- Smith, Adam 66
- Socionomic Institute 158
- Socionomics 130, 157
- Soros, George 312
- Spinning top candlestick 300
- Spread of market 58
- Standard deviation 58–60
- Standing orders 83
- Starbucks Index 337
- Stock broker case study 386–7
- Stock fundamentals 342

- Stoller Average Range Channels (STARC)
 - Bans 236, 237
- Stop hunting 83
- Stop-loss order 82–4
- Strong hands 112
- Structural fractal 300
- Style diversification 318–19
- Supply curve 45–7
- Supply and demand 45–8
- Supply/demand index 303, 305, 397–403
- Symmetrical triangles 88–91
- Systematic market trading 14

- Tape readers 199
- Tech bubble 39
- Technical analyst case study 380–2
- Technical market analysis 3–5, 39, 162
- Technomonic analysis 390–4
- Tesco weekly – classic rectangle 93–5, 344–5
- Theta 254
 - effect 328–9, 331
 - short 325
- Time decay, *see* Theta
- Time fractals 395–7
- Time value effect 331
- Trader vs investor 309–10
- Trading pits 68–70
- Trading regime analysis 310
 - for the long term 16–17
- Trading regime analyst 18
- Trading Regime Indicator (TRI) 64, 191, 291–5
- Trading regimes 14–16
- Traditional technical analysis techniques 16
- Trend, definition 15–16
- Trend exhaustion 16, 63–5
- Trend following 5, 331–6
 - diversification 332–3
 - long term 332
 - models 239
 - strategies 9–13, 14
 - volatility breakdown 334–6

- Trend Following Performance Indicator (TFPI) 281–9, 403
- Trending regime, psychology of 73–4
- Trend lines 5, 407–9
- Trend mode 15
- Triangles 80–91
- ascending 81–6
 - descending 86–8
 - expanding 101–3
 - symmetrical 88–91
- True Range 194, 231, 232
- Turtle Type System 225–6, 227–30, 234, 276, 291, 352
- Uptrend resistance 413–14
- US 10-year Yield Weekly – doji trading range 135–6
- USDCAD monthly – doji trading range 134–5
- USDCAD weekly 158–9
- and Bollinger Band Width 307
 - and RSI 214–16
- USDCHF
- monthly – Turtle type system 229–30
 - weekly – Turtle type system 228–9
- US Department Stores sector 350
- and MA Envelopes 351
- USDJPY
- daily – A–B–C range trade 173–5
 - daily and STARC Bands 236–7
 - weekly – classic descending triangle 87–8
- USDMXN monthly – falling wedge 126–7
- US Dollar Index
- Structural Supply Demand Index 399–400
 - Supply/Demand Index 400–2
- US Dollar Index future (DXY) 300
- US Dollar Index monthly 302
- US Dollar Index weekly 148
- classic head and shoulder 105–9
 - standard deviation 62–3
- US Dollar spot index 300
- US Gross Domestic Product – Dow Jones Industrials 345–6
- US Industrial Production Index and MA Envelopes 360–1
- US Investment Bank Sector weekly 118–19
- US Long Bond 47
- US Pharmaceuticals
- relative daily 355–6
 - relative monthly and MA Envelopes 353–5
- US Treasury
- 10-year Yield and ADX 195–7
 - 10-year Yield weekly 292–4
 - 2-year Yield monthly – rate of change 211–12
 - 30-year Bond future 70–1
- Value investing 336–41
- Value strategy 14
- Value style of investment 318
- Variable P&L 314
- VIX Index 44, 54
- Volatility 13, 15, 42, 43–76, 356
- defined 48–50
 - forecasting 54
 - historical, calculation 50–3
 - implied 55–7
 - probability of 18
 - vs price expansion and contraction 60–3
- Volatility analysis 17, 310
- Volatility breakout trade 303, 304, 336
- Volatility conditions 18, 319, 320
- Volatility smile 261–74
- Volatility Smile Index (VSI) 266–7
- Volume 143–59
- analysis 89
 - considerations 303
 - increasing on breakouts 144–5
 - spikes after trending moves 145–50
- Vol of vol 65–6
- V-shaped recovery 110–11

Wave Principle 173
Weak hands 112
Wedges 119–27
 falling 124–7
 rising 119–24
Weekly rule 227
Williams %R 219–23, 302

World Cup of Trading Championship
 219–20

Zell, Sam 32–3
Zone of extreme excitement 255–6
Zone of extreme relaxation
 255

Index compiled by Annette Musker