



Durham E-Theses

On the profitability of technical trading

WATSON, TOBY,DANIEL

How to cite:

WATSON, TOBY,DANIEL (2009) *On the profitability of technical trading*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/198/>

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

On the profitability of technical trading

Toby Daniel Watson

Submitted for the degree of Doctor of Philosophy
September, 2009

Abstract

The sole use of price and related summary statistics in a technical trading strategy is an anathema to weak-form market efficiency. In practice, however, traders actively use technical analysis to make investment decisions which makes this an important, but often neglected, area for study. This thesis includes four empirical chapters, which provide important evidence on the profitability of technical trading. The results from the detailed analysis undertaken in this thesis have broad relevance to both academics and those in the investment community.

Existing research has been predominantly confined to evaluating basic technical trading rules, such as moving averages. Crucially, this ignores chart patterns. Widely employed by practitioners, such patterns form a vital part of technical analysis. As the most important price pattern, the head and shoulders pattern is subjected to detailed and thorough examination in this thesis. A significant contribution is made by evaluating formations recognised and used by traders, in sharp contrast to limited existing studies. Furthermore, a new method is developed to establish how quickly profits from a head and shoulders strategy decay, which has important implications for traders.

Existing research has identified both reversal and relative strength effects in financial asset returns. A key separator between these two findings is the formation and holding time over which portfolios of winners and losers are evaluated. Motivated by this, a very large sample of ultra high-frequency data is used to investigate intraday momentum and reversal effects. As well as being an important contribution to research in this field, the results are, once again, of relevance to practitioners.

The need for further research into technical analysis is clearly demonstrated by point and figure charting. Whilst traders have made consistent use of the technique for around a century, the amount of existing research is extremely small. Point and figure has attractive data filtering properties, clear trading rules and is particularly suited to intraday technical analysis. Again, using a very large sample of high-frequency data, a detailed evaluation of the profitability of a point and figure trading strategy is undertaken.

On the profitability of technical trading

Toby Daniel Watson

A Thesis presented for the degree of
Doctor of Philosophy



Department of Economics and Finance
Durham Business School
Durham University
September, 2009

Contents

Abstract	i
List of Figures	iv
List of Tables	vi
Declaration	ix
Acknowledgements	x
1 Introduction	1
2 Advanced Technical Analysis: The head and shoulders pattern	8
2.1 Introduction	9
2.2 Organisation	13
2.3 A review of the literature	14
2.3.1 Introduction to the issues	15
2.3.2 Early academic research	17
2.3.3 Basic technical analysis implementations	18
2.3.4 Advanced technical analysis implementations	23
2.3.5 Pattern recognition and the head and shoulders	25
2.3.6 Technical analysis in the markets	34
2.3.7 Literature review conclusions	38
2.4 Data and methodology	40
2.4.1 Nature and breadth of data	40
2.4.2 Identifying peaks and troughs in price data	41
2.4.3 Kernel smoothing	42
2.4.4 Detecting head and shoulders patterns	54
2.4.5 Recognising head and shoulders patterns	61
2.4.6 Evaluating head and shoulders patterns	64
2.4.7 Conclusions	67
2.5 Empirical Results	68
2.5.1 Summary Statistics	68
2.5.2 Is the head and shoulders profitable?	70
2.5.3 Transaction and short selling costs	80
2.5.4 Bootstrap tests	82
2.6 Conclusions	86

3	Removing the Straight-Jacket on Technical Analysis	88
3.1	Introduction	89
3.2	Organisation	91
3.3	Practitioner literature	92
3.3.1	The history of technical analysis	92
3.3.2	The head and shoulders pattern	95
3.4	Data and methodology	97
3.4.1	Nature and breadth of data	97
3.4.2	Identifying peaks and troughs in price data	97
3.4.3	Identifying head and shoulders patterns	98
3.4.4	Conclusions	102
3.5	Empirical results	103
3.5.1	Transaction and short selling costs	112
3.5.2	Bootstrap tests	113
3.6	Conclusions	116
4	Intraday Reversal or Relative Strength?	119
4.1	Introduction	120
4.2	Literature	124
4.3	Data and methodology	132
4.4	The Returns of Intraday Winner and Loser Portfolios	136
4.4.1	Subperiod analysis	141
4.4.2	Day of week and intraday relative strength portfolios	144
4.4.3	Time of day and intraday relative strength portfolios	145
4.4.4	Size and intraday relative strength portfolios	150
4.5	Conclusions	152
5	Point and Figure Trading	157
5.1	Introduction	158
5.2	Point and Figure Analysis and Literature	161
5.2.1	Construction of point and figure charts	161
5.2.2	Point and figure literature	169
5.3	Data and methodology	174
5.4	The Profitability of Point and Figure Trading Rules	183
5.4.1	3-box chart construction	184
5.4.2	5-box chart construction	190
5.5	Conclusions	193
6	Conclusions	195
6.1	Thesis overview and importance of findings	196
	Bibliography	200

List of Figures

2.1	The head and shoulders top pattern	27
2.2	The head and shoulders bottom pattern	28
2.3	Formation of head and shoulders patterns	42
2.4	Examples of the application of various kernels.	47
2.5	Bandwidth selection	51
2.6	Bandwidth selection	52
2.7	Bandwidth selection	53
2.8	Bandwidth selection - globally optimised	55
2.9	Bandwidth selection - locally optimised	56
2.10	Determining peaks and troughs with the use of the signum function. $h = 2$	57
2.11	Head and shoulders pattern detection using the rolling windows .	59
2.12	Diagram of a head and shoulders patterns with extrema labelled .	62
2.13	Trading the head and shoulders pattern	63
2.14	Cumulative mean standard deviations over 500 simulations.	83
3.1	Diagram of an artificial head and shoulders pattern	100
3.2	Cumulative mean standard deviations over 500 simulations.	114
4.1	Returns of relative strength portfolios	140
4.2	Buy-sell portfolio returns by calendar month	143
4.3	Buy-sell portfolio returns by day	146
4.4	Returns of Relative Strength Portfolios by hour	149
4.5	Buy-sell portfolio returns by market capitalisation	153
5.1	Standard line chart of stock price	161
5.2	Point and figure chart of the same series as Figure 5.1, with a box size of \$1 and reversal of 3 boxes	162
5.3	Intraday price chart of Microsoft Corporation - 3 January, 2005	166
5.4	Point and figure chart of stock price shown in Figure 5.3 with a box size of \$0.02 and reversal of 3 boxes	166
5.5	Point and figure chart of stock price shown in Figure 5.3 with a box size of \$0.02 and reversal of 5 boxes	167
5.6	Point and figure chart of stock price shown in Figure 5.3 with a box size of \$0.04 and reversal of 3 boxes	168
5.7	Point and figure chart of stock price shown in Figure 5.3 with a box size of \$0.04 and reversal of 5 boxes	169
5.8	Double Top buy signal (B1)	176

5.9	Double Bottom sell signal (S1)	177
5.10	Double Top buy signal with rising bottom (B2)	178
5.11	Double Bottom sell signal with declining bottom (S2)	178
5.12	Breakout of triple top (B3)	179
5.13	Breakout of triple bottom (S3)	179
5.14	Ascending Triple Top (B4)	180
5.15	Descending Triple Bottom (S4)	180

List of Tables

2.1	Key features defining head-and-shoulders patterns, applied to the idealised pattern	64
2.2	Frequency counts for occurrences head and shoulders (HS) and inverse head and shoulders (IHS) patterns in UK stock data	69
2.3	Head & Shoulders returns for 1980-2003 with extrema identified using a globally optimised bandwidth	71
2.4	Head & Shoulders returns for 1980-2003 with extrema detected using a 30% of the globally optimised bandwidth ($h \times 0.3$).	74
2.5	Head & Shoulders returns for 1980-2003 with extrema detected under global bandwidth optimisation with $h \times 2$	75
2.6	Head & Shoulders returns for 1980-2003 with extrema detected using a locally optimised bandwidth.	77
2.7	Head & Shoulders returns for 1980-2003 with a trade lag of ≤ 5 and extrema detected using a locally optimised bandwidth.	78
2.8	Head & Shoulders returns for 1980-2003 with a trade lag of > 5 and ≤ 10 and extrema detected using a locally optimised bandwidth.	80
2.9	Bootstrap results from 500 simulated series compared to the actual price series.	81
3.1	Key features defining head-and-shoulders patterns	95
3.2	Head & Shoulders returns for 1980-2003 under pattern specification HSB and IHSB.	104
3.3	Head & Shoulders returns for 1980-2003 under pattern specification HSB and IHSB with a trade lag of ≤ 5	106
3.4	Head & Shoulders returns for 1980-2003 under pattern specification HSB and IHSB with a formation period of 65 days.	107
3.5	Head & Shoulders returns for 1980-2003 under pattern specification HSB and IHSB with a formation period of 65 days and a trade lag of ≤ 5 days.	108
3.6	Head & Shoulders returns for 1980-2003 under pattern specification HSC and IHSC.	109
3.7	Head & Shoulders returns for 1980-2003 under pattern specification HSC and IHSC with a trade lag of ≤ 5 and pattern formation period of 35 days.	110
3.8	Head & Shoulders returns for 1980-2003 under pattern specification HSC and IHSC with a pattern formation period of 65 days.	111

3.9	Head & Shoulders returns for 1980-2003 under pattern specification HSC and IHSC with a trade lag of ≤ 5 and a formation period of 65 days	112
3.10	Bootstrap results from 500 simulated series compared to the actual price series.	113
4.1	Descriptive Statistics	135
4.2	Returns of relative strength portfolios	137
4.3	Returns of relative strength portfolios by calendar month	142
4.4	Returns of relative strength portfolios by day of week	144
4.5	Returns of Relative Strength Portfolios by Hour	147
4.6	Returns of relative strength portfolios by market capitalisation	151
5.1	Returns of Point and Figure Trading Rules, where $R = 3$	185
5.2	Returns of Point and Figure Trading Rules, where $R = 5$	191

Declaration

No part of this thesis has been submitted elsewhere for any other degree or qualification in this or any other university. It is all my own work unless referenced to the contrary in the text.

Copyright © 2009 by TOBY WATSON.

“The copyright of this thesis rests with the author. No quotations from it should be published without the author’s prior written consent and information derived from it should be acknowledged.”

Acknowledgements

I would like to thank Professor Phil Holmes and Professor Krishna Paudyal for their support throughout this research. Their advice on the direction of the work and helpful comments, together with encouragement, has proved invaluable. In addition, Professor Robert Sollis of Newcastle Business School provided many helpful comments during the early stages of this research. Discussions with members of the Department of Economics and Finance, part of Durham Business School, have also proved extremely helpful. In particular, informal conversations with Professor John Ashworth and Professor Simon Parker have been of great assistance. My parents, Annie and Michael, deserve sincere thanks for their support of my studies. I am indebted to my wife, Sarah, for her support and patience throughout.

Dedicated to
Sarah

Chapter 1

Introduction

“It seems very clear that under scientific scrutiny chartreading must share a pedestal with alchemy. There has been a remarkable uniformity in the conclusions of studies done on all forms of technical analysis. Not one has consistently outperformed the placebo of a buy-and-hold strategy. Technical methods cannot be used to make useful investment strategies. This is the fundamental conclusion of the random-walk theory.”

Malkiel (1999)

Technical analysis involves the sole use of price and related summary statistics, such as volume, to inform trading decisions. Given its long-standing use in financial markets, technical analysis has naturally become a focus of academic study. In part, this is because profits accruing from a strategy constructed entirely around the analysis of past prices runs counter to the least restrictive form of market efficiency. The above quotation from Malkiel (1999) expresses this opinion, based on a belief in efficient markets. This thesis examines several important areas of technical analysis and finds that there is strong empirical evidence that opposes this point of view.

Forecasting future price changes of financial assets with the aid of charts of past prices has a long history of use by investors and traders. For example, Nison (1994) describes the development of candlestick charts, which provide a visual representation of the opening, closing, high and low prices for a discrete period. It is shown that such charts may have been used as early as the 1700s by traders in what was, in effect, the first rice futures market in Japan. Furthermore, it is certain that traders plotted candlestick charts and used them to inform trading decisions by the late 1800s. The so-called “book method”, which was an early version of point and figure charting, was also in active use by 1900. Indeed, Charles Dow published a *Wall Street Journal* editorial on the subject in 1901 (Murphy, 1999). Thousands of books on technical analysis aimed at traders have since been published, with many different forms of technical trading strategies proposed, to be employed across the whole gamut of financial markets, including equities, foreign exchange and futures. Indeed, all professional trading platforms, such as Reuters and Bloomberg, can

perform technical analysis. The common thread is the sole use of past price data for making buy and sell decisions.¹

Importantly, it is clear that the continuing non-academic interest in technical analysis translates into active use in the markets. For example, Taylor and Allen (1992) conducted a survey of foreign exchange traders in London. The results showed that where respondents employed in-house technical analysts, there was a greater tendency for them to initiate trades as opposed to in-house economists. Other surveys also provide convincing evidence that traders make significant use of technical analysis, either in isolation or in conjunction with fundamental analysis (for example, Lui and Mole, 1998; Gehrig and Menkhoff, 2006; Cheung et al., 2004). If technical trading strategies do not provide economically valuable information, then their continuing use proves somewhat perplexing, and provides strong motivation for increased academic study.

Technical analysis covers a multitude of different techniques and strategies to utilise price data. For example, moving averages, relative strength, trend indicators and price patterns. There are also innumerable chart styles, such as bar charts, candlestick charts and point and figure charts. However, previous academic research in this area has largely concentrated on what can be termed 'basic' technical analysis, such as moving averages. This is partly because it is relatively easy to construct algorithms to evaluate the profitability of basic technical trading strategies.

However, we know considerably less about the profitability of what I term, for the purposes of this thesis, 'advanced' technical trading strategies. Advanced technical trading strategies are generally concerned with detecting and evaluating visual patterns displayed on charts of past price data. Whilst formations approximating a particular specification are usually clear to the human eye, it is a considerable problem to develop algorithms to allow the evaluation of advanced technical analysis by computer. Price patterns are therefore very different from

¹Section 2.3.1 provides a detailed examination of the nature of technical analysis and the range of trading strategies it encompasses.

trading strategies such as the moving average, where buy and sell signals can be easily derived from a vector of past prices. It is only comparatively recently that appropriate econometric methods and sufficient computational power has existed to allow a full investigation of advanced technical analysis.

A related point concerns high-frequency data. Many technical trading methodologies are agnostic of the time frame over which they can be applied—for example, being equally valid using weekly charts, daily charts and intraday charts. Thus, a 50-period moving average could be employed over 50 weeks, 50 days or 50 minutes. In addition, many other strategies are specifically proposed as being useful over short time horizons. Investigating technical trading strategies employed using high frequency data, as I do here, is particularly important given the increasing numbers of day traders. Professional traders, and hedge funds in particular, also employ program trading strategies that utilise technical analysis. Yet we still do not know very much about the profitability of such strategies.

High-frequency data has been available for some time from sources such as the New York Stock Exchange. However, it is only relatively recently that such data has been readily available to the academic community and, again, that computational power has allowed researchers to take full advantage of this. It is now possible and increasingly pressing that we investigate the profitability of technical trading strategies with high-frequency data.

In addition, most existing research has not succeeded in evaluating and applying technical trading strategies as they are actually employed by traders, when making buy and sell decisions. For example, there is often a clear disparity between the head and shoulders pattern that is consistently seen in the literature aimed at practitioners and that which is evaluated in academic research. This is partly because of the aforementioned problem of computational power and suitable methodology. This study, however, makes considerable progress in addressing this issue.

Given its long history, technical analysis has seen the development of innumerable indicators, patterns, chart types and trading strategies. Partly due to the depth and breadth of the subject, academic investigation has been severely limited or even non-existent into a great many aspects of technical trading.

The sparsity of empirical evaluation into areas of technical analysis, the lack of investigation into many trading strategies at time horizons employed by traders, and the scant knowledge about the profitability of advanced technical trading strategies makes technical analysis a compelling and timely area for study. This thesis seeks to examine the nature and profitability of a number of important technical trading strategies, and make a significant contribution in several important areas.

First, advanced technical analysis is investigated through an evaluation of the head and shoulders pattern. This is the most prominent price pattern in technical analysis, and is exhibited in most practitioner—as opposed to academic—texts. Chapter 2 evaluates the limited existing research on the head and shoulders pattern, and undertakes an empirical investigation based on a large sample of UK securities. This research makes several important contributions. Specifically, existing research does not evaluate the head and shoulders pattern for profitability as part of a trading strategy. In this work, the returns to a head and shoulders based trading strategy are evaluated. Furthermore, the introduction of a ‘trade lag’ allows investigation of the speed of decay of profits owing to the detection of head and shoulders patterns. The chapter also looks at the formation of head and shoulders patterns over a variety of time horizons, ranging from 1 to 60 days, and also investigates four subperiods, providing an insight into whether the performance of the strategy is conditional upon the state of the market.

Second, given the importance of building a greater understanding of the performance of advanced technical strategies, Chapter 3 significantly extends the analysis of head and shoulders patterns in a number of valuable ways. Firstly,

one of the problems in much of the existing research is that the pattern geometry of head and shoulders patterns studied do not closely correspond to those used by traders. Accordingly, a close examination of the practitioner literature is undertaken—all too often ignored in academic work on technical analysis—and the geometric specification of patterns is evaluated in detail. Most importantly, new specifications are developed which are more closely aligned to what would be utilised by practitioners. It is determined that one of the critical features in any study of advanced technical analysis is the detection of peaks and troughs in past price data, which serve as the building blocks for recognising price patterns. Given the importance of finding local maxima and minima in noisy price data, a further contribution of this chapter is to evaluate an alternate method of their detection.

There are clear links between technical trading and the existing literature on momentum and reversal in stock returns. As well as being solely concerned with past price (and associated return) data, ignoring ‘traditional’ measures of valuation such as discounted cash flows, there is an important relationship to trends in technical analysis. Murphy (1999, p.49) states that “the concept of a *trend* is absolutely essential to the technical approach to market analysis”, and broadly defines a trend as “simply the direction of the market.” Of course, there are different lengths of trend lasting from just a few hours to many years. It is therefore possible to consider that momentum or reversal effects must be clearly present for technical analysis to succeed.

Following, in particular, Jegadeesh and Titman (1993) and De Bondt and Thaler (1985), a large literature has developed on reversal and momentum in stock prices and returns. The time horizon over which the performance of past winners and past losers is evaluated has been crucially important in bridging the gap between the strands of momentum and reversal literature. Given the importance of this factor, it is somewhat surprising that we lack knowledge about momentum and reversal effects at the intraday level. As trends and price momentum and reversal

are at the very core of technical trading, then this is an important area of research if we are to establish the profitability of intraday technical trading.

Chapter 4 investigates short-term momentum and reversal strategies. This chapter, in contrast to the limited existing research in this area, uses an extremely large sample of high-frequency trade data from the New York Stock Exchange Trade and Quote (TAQ) database to investigate short-term momentum and reversal effects.

Having recognised that technical analysis is an extremely broad subject area, we still know next to nothing about some chart types and trading strategies. One under-investigated area is point and figure charting. Point and figure charting uses a filtering method to plot price changes on a chart that differs in nature from the conventional line and bar charts that are most commonly seen. As well as considering the lack of research into the technique, it forms a compelling topic for study for two main reasons: it has a very long history of use by traders and is still in active use today; furthermore, there are attractive properties of the technique in filtering noisy price data, and the charts produced lend themselves to recognising patterns computationally.

Empirical work in Chapter 5 investigates this compelling area, and makes use of a large sample of high-frequency trade data. This provides a crucial element of the analysis given that point and figure charting was originally intended to be used by floor traders with tick data.

Together, these four empirical chapters provide new insight into a number of important aspects of technical trading. The findings are especially relevant, and of interest to those outside the academic community, given the continued use and weight accorded to technical analysis by market practitioners. Chapter 6 provides a summary of findings.

Chapter 2

Advanced Technical Analysis: The head and shoulders pattern

2.1 Introduction

The existing body of research in technical analysis is mainly concerned with looking at simple trading rules that do not accurately capture the activity of professional traders. Traders often use visually complex chart patterns in price data to inform their decisions in place of, or in combination with, basic indicators such as moving averages. I term this ‘advanced technical analysis’, and propose this definition for the move beyond simple strategies such as moving averages, towards a recognition of these predominantly ‘visual’ patterns in price data. This is the first study that rigorously examines the profitability of a trading strategy based on advanced technical analysis, using the head and shoulders pattern.¹ Several innovations give rise to a major contribution to the existing literature. Most notably, by developing the completely new idea of a ‘trade lag’, it is possible to evaluate how quickly any profits from head and shoulders patterns are arbitrated away. Furthermore, evaluation of head and shoulders profitability over a number of different time horizons ranging from 1 to 60 days allows the persistence of profits to be established. Little is known about this based on current research. The study is supported by a large dataset for UK stocks running from January 1, 1980 to December 31, 2003. The core research question addressed in this study is to what extent head and shoulders patterns lead to a profitable trading strategy, in the context of a large sample of UK stocks.

Technical analysis has considerable pedigree in the financial markets. Brock et al. (1992, p.1731) point out that “[it] is considered by many to be the original form of investment analysis, dating back to the 1800s”. Technical analysis retains an important role in the financial markets with all major investment banks employing dedicated staff—if not whole departments—to examine patterns and trends in

¹Whilst a limited number of studies have made use of the head and shoulders pattern, there are many shortcomings which will be discussed in more detail below. Comprehensively addressing these limitations is one of the aims of this work.

past prices.^{2,3} The long-established use of past-price history in making investment decisions, together with any availability of abnormal profits from technical trading strategies running counter to weak-form market efficiency, is seemingly difficult to reconcile with studies showing that profits from technical trading strategies appear to persist. This gives clear motivation for the study of technical analysis: why does the use of technical analysis persist, and is the lack of an answer to this question in part because existing research has largely ignored the type of technical analysis actually practised by market participants?

The head and shoulders pattern is one of the most prominent and long-standing chart patterns and regarded as one of the most informative by traders. For example, Achelis (2001, p.246) describes it as “the most reliable and well-known chart pattern,” and Murphy (1999, p.103) determines that the head and shoulders as “probably the best known and most reliable of all major reversal patterns”. The head and shoulders pattern can therefore be considered to be the best example of advanced technical analysis. Consequently, it is selected as the basis for this work. There is a long history of the head and shoulders pattern being used by technical analysts; for example, Edwards and Magee (1948) identified the importance of head and shoulders patterns in stock price charts. Such a long history of active use of the pattern negates claims of data mining.

The central motivating factor of this work is based upon examining the hitherto under-investigated subject of advanced technical analysis. Whilst traders have been using strategies employing pattern recognition for a long time, this has not been a prominent feature of academic research.⁴ The apparent lack of interest in technical analysis in the literature is partly down to the computational power required to systematically evaluate complex technical analysis, replicating what traders use the

²Of course, the use of technical analysis is heterogeneous across bank functions.

³For example, see Taylor and Allen (1992) for details and survey evidence on the use of technical analysis by traders; more details of the use of technical analysis in financial markets today can be found in Section 2.3.6 on page 34.

⁴These claims will be supported by the literature review undertaken below (Section 2.3 on page 14).

human eye for, not being available until relatively recently. Harnessing this power with the use of an algorithm to recognise head and shoulders patterns allows this study to critically evaluate the profitability of these formations in the context of a trading strategy. The landmark study by Lo et al. (2000) rekindled academic interest in technical analysis by looking at patterns in price data. Lo et al. do not, however, evaluate the profitability of trading strategies based on such patterns. Instead, the difference in unconditional 1-day returns versus 1-day returns conditioned on the existence of patterns is evaluated as a proxy for patterns' informational content. As such, there is a major shortcoming in our knowledge—we do not know whether such patterns are actually *useful* in an economic sense. Furthermore, the study only looks at 50 stocks per period under investigation, using only a small sample of US data.⁵

By testing whether price patterns contain information that can be employed profitably in a trading strategy, this study addresses the key shortcoming of Lo et al., and significantly extends our knowledge of technical analysis. Furthermore, the concept of the 'trade lag' is developed. This new approach allows an investigation of how quickly any profitability associated with trading on head and shoulders patterns is arbitrated away. This is achieved by imposing a variety of different restrictions on the elapsed time between the detection of a head and shoulders pattern and a trade occurring. Furthermore, to investigate the profitability of head and shoulders patterns within a trading strategy, a number of different trade horizons from 1 to 60 days are evaluated. In addition, a larger dataset is employed providing a high degree of robustness to results. Taken together, this study therefore provides a significant and original contribution over and above the results and conclusions obtained in previous work.

The major contributions of this study can be enumerated as follows:

⁵Dawson and Steeley (2003) replicate Lo. et. al.'s methodology for the UK, but are subject to the same shortcomings. Savin et al. (2007) provide a very recent study looking at patterns in a large sample of US data. However, this work is subject to important limitations which will be discussed below, and fully addressed by this study.

1. The primary and most significant contribution is to the developing literature on 'visual' technical analysis patterns. Unlike previous research (in particular Lo et al. (2000) who also look at the head and shoulders pattern), this study seeks to evaluate whether the head and shoulders is actually profitable for traders. This is important as this form of technical analysis is actively used in the markets.
2. A large dataset of UK equity data is utilised. All daily stock price data is collected for the period January 1, 1980 to December 31, 2003; this allows the portfolio of the 350 largest stocks by market capitalisation to be resampled annually. Head and shoulders patterns seem to be more prevalent in larger stocks. Greater liquidity in larger stocks increases the likelihood of head and shoulders patterns occurring.⁶ As will be seen below, much of the existing research (for technical analysis in general) is concentrated on a small number of currency pairs, a limited sample of stocks or index data. The length of the sample encompasses a range of market conditions.
3. Little is known about the persistence of head and shoulders profits after their formation. This study addresses this question by scrutinising profitability for a variety of different holding periods ranging from 1 day to 60 days.
4. Crucially, this study introduces an entirely new concept termed the 'trade lag'. This allows an evaluation of how quickly any profits from head and shoulders patterns are arbitrated away. In other words, do 'fresher' patterns perform better?⁷

⁶One reason is that in smaller illiquid stocks larger price changes may lead to patterns lacking 'symmetry' and therefore being unrecognisable. Indeed, the pattern specifications outlined below included symmetry criteria. Supporting this point from a practitioner perspective, Bulkowski (2005, p.4) excludes stocks that 'did not have consistently large daily price ranges (too thinly traded or volatile).' Indeed, the practitioner literature mostly draws examples of the occurrence of chart patterns from larger stocks. Furthermore, there is a link between technical analysis and liquidity; for instance, Kavajecz and Odders-White (2004) show that support and resistance levels are related to the depth of the order book. Order book depth is greater in larger stocks.

⁷By definition, a local maximum or minimum is only known after its formation; the trade lag

The results and conclusions will not only advance the debate on technical analysis, but will be useful to several groups. Given their extensive use of technical analysis, traders will be keen to know if price patterns can indeed generate superior returns. Investors are also highly interested in this topic in the context of the debate between technical and fundamental analysis. In summary, this study forms the vital next step in evaluating profits from advanced technical analysis, focussing on the head and shoulders pattern.

2.2 Organisation

In addressing the core research question, this chapter is divided into three further sections. First, a critical review of the literature is undertaken with a view to demonstrating the gaps that motivate this work and allow the framing of the central research questions (Section 2.3). Second, the data and methodology section details the dataset used and the methodology adopted (Section 2.4). In this section, the key steps necessary for the detection of head and shoulders patterns are identified: detecting peaks and troughs in noisy price data and establishing the geometric properties of the head and shoulders pattern.

Empirical results (Section 2.5) are presented and discussed with a view to addressing the issue of the profitability of the head and shoulders pattern. Tables are presented to show the returns contingent on the trade lag as well as over a number of trade time horizons. Conclusions to the study are presented in Section 2.6.

allows evaluation of whether patterns formed more recently in relation to the current time period perform better.

2.3 A review of the literature

“There is no way of making an expected profit by extrapolating past changes in the futures price, by chart or any other esoteric devices of magic or mathematics. The market quotation already contains in itself all that can be known about the future and in that sense has discounted future contingencies as much as is humanly possible.”

(Samuelson, 1965, p.44)

Samuelson succinctly expresses the opinion that in an efficient market we would not expect to be able to make profits through technical analysis. This review of the literature shows the increasing interest in technical analysis by researchers, often demonstrating that profits can be shown, in contradiction of weak-form efficiency. Existing research is classified accordingly into two broad groups: First, basic studies of technical analysis, which are recognisable by the evaluation of simple rules and trading strategies such as filter rules and moving average crossovers. Second, ‘new’ studies of technical analysis. This more recent work tends to possess more robust econometric methodology. More advanced technical analysis strategies—including pattern recognition—are also included in this group. With the exception of advanced technical analysis, pattern recognition and studies concerning the head and shoulders pattern itself, this review is not intended to be exhaustive.^{8,9} Rather, its strength is in pointing the reader to the papers and research that has shaped academic understanding of technical analysis. Before this, however, it is important to establish a firm grip on what constitutes technical analysis, and this is addressed in the next section.

⁸A large body of work investigates simple technical strategies; however, as noted, it is the complex and predominantly visual patterns that are of specific interest here.

⁹Park and Irwin (2007) provide a useful general overview of the literature in the area of technical analysis.

2.3.1 Introduction to the issues

Technical analysis (or chartism as it is often referred to by investment professionals) is an “attempt to forecast prices by the study of past prices and a few other related summary statistics about security trading” (Brock et al., 1992, p.1731). This indicates the reason that it has often been held in such disdain by academics; in focusing on past prices alone, technical analysis directly contradicts weak-form market efficiency, which states that it should not be possible to earn excess returns from studying past price movements. Technical analysts (‘technicians’ or ‘chartists’) have created many ways to use historical prices in an attempt to extrapolate future movements, ranging from basic averaging indicators to visually oriented chart patterns which are considerably more difficult to express algebraically in the context of academic study. Achelis (2001) and Bulkowski (2005) show just how many technical indicators, patterns and strategies have been created and employed by technical analysts.

Earlier studies of technical analysis generally provided support for weak-form market efficiency and determined that a range of basic indicators did not generate abnormal returns (Fama and Blume, 1966; Jensen and Benington, 1970). However, there has recently been renewed interest in examining a broad range of technical indicators and strategies, which has developed largely in tandem with the discovery of various ‘anomalies’, such as day of the week effects. In addition, fundamental investment strategies have produced more evidence against semi-strong form market efficiency, for example contrarian value investment (Lakonishok et al., 1994; La Porta et al., 1997; Fama and French, 1998). It should, however, be made clear that in the strictest sense technical analysts are only concerned with past prices and related summary statistics. Related summary statistics essentially only refers to volume and open interest.

Several points should be considered when reviewing the body of literature. First, until relatively recently, the lack of available computational power imposed a

restriction on the study of technical analysis. When scholars first became interested in technical analysis it was too ‘computationally expensive’ to test even basic technical trading rules (such as the moving average) on large datasets. Of course, this problem is particularly acute for high-frequency intraday data.¹⁰ Thus, much early work focusses upon the past values of market indices and, in particular, the Dow Jones Industrial Average. It was also impossible to investigate advanced technical analysis strategies. Recognition of patterns in price data is very computationally intensive.¹¹

Furthermore, detecting chart patterns in price data crucially depends on having a reliable method to extract useful maxima and minima points. A number of possibilities exist; for example, a moving window or the methodology introduced by Bry and Boschan (1971) to identify turning points in the business cycle. Latterly, smoothing methods such as kernel regression have become a popular technique to isolate key points in a series dominated by noise and volatility. However, when applied to large datasets these approaches are all computationally intensive, which is one explanation for a lack of thorough investigation of price patterns until relatively recently.

We should also not ignore advances in econometric methodology. For example, greater recognition of non-stationarity in financial time series and the presence of time-varying returns may have invalidated aspects of much of the earlier research. There have been huge improvements in the ways that we can analyse the results of technical analysis strategies by treating them as forecasting models.

¹⁰Later chapters in this thesis make use of increasingly cheap computational power to examine technical analysis in a new light, using ultra high-frequency data from the NYSE Trade and Quote database.

¹¹For instance, even with modern computers and the highly optimised algorithms I have developed for the empirical work in this thesis, the bootstrap results presented later in this chapter require over 2,000 computational hours for 500 simulations using a 2007 Intel® Xeon® workstation-class processor. This equates to around 83 days when run in sequence.

2.3.2 Early academic research

We could continue to look at the prelude to the formation of the efficient markets hypothesis, for example, refutations by Alexander (1961), Alexander (1964) and Weintraub (1963). However, while for obvious reasons technical analysis has been a victim of research supportive of efficient markets, the focus of this study is advanced technical analysis. This section presents some of the early research on technical analysis that appears in the literature, with the purpose of demonstrating and evaluating key concepts that are a necessary building block for this chapter and the rest of the thesis.

Scholarly interest in technical analysis can be traced as far back as Cowles (1933), who undertook an examination of stock price forecasting methods. This included looking at technical trading and, in particular, the activities of William Hamilton in employing Dow Theory (Hamilton, 1922). Dow Theory was developed by Charles Dow, the editor of the Wall Street Journal in the late 1800s. Underpinning his ideas was the concept that the market moved in trends, with minor and medium trends being able to occur in the opposite direction to the main trend. The most interesting proposition was that of an 'accumulation phase', where informed investors traded against the market at the start of the movement, and sold towards the end of the trend in a 'distribution phase'. In the distribution phase, informed investors were thus taking profits as new and less informed individuals belatedly bought. However, Dow also concluded that the market quickly impounded new information when released, which would now be seen in the context of the EMH. However, Dow was not especially interested in proposing a trading strategy based on these ideas. Indeed, it was later Wall Street Journal editors, including Hamilton, who developed his work and coined the expression Dow Theory. The early scholarly research by Cowles discovered that so called 'stock market forecasters' concerned with pursuing these early theories were, in fact, not particularly successful in their forecasts. However, the forecasts obtained by Hamilton in employing Dow Theory

were not impressive, and certainly insufficient to comprise a profitable trading strategy.

Roberts (1959) initiates what could be viewed as the ‘classical’ view of technical analysis as an anathema to serious scholars. Building on emerging research into the random walk nature of prices—particularly Kendall (1953)—Roberts asserts that price patterns are merely an insignificant artefact in price data. Firmly refuting technical analysis, it is stated that:

“In light of this intense interest in patterns and of the publicity given to statistics in recent years, it seems curious that there has not been widespread recognition among financial analysts that the patterns of technical analysis may be little, if anything, more than a statistical artifact.”

(Roberts, 1959, p.1)

However, when referring to ‘price patterns’, Roberts does not speak of chart patterns, such as the head and shoulders, as we would now view them. Instead, he is mainly concerned with pointing out the random nature of the Dow Jones index and thus casting doubt on the work of some financial analysts. Even so, the view of technical analysis presented was representative of the academic perspective at the time and indeed to some extent today. It is, however, not until the 1960s that scholarly investigation began to fully investigate aspects of technical analysis.

2.3.3 Basic technical analysis implementations

As noted above, whilst practising technical analysts may regard simple technical strategies as somewhat elementary, their investigation makes up the bulk of literature.¹² Having already discussed some of the perceived shortcomings in looking

¹²Note that when referring to practising technical analysts or ‘technicians’ this means investment professionals who utilise these strategies.

solely at these strategies, it is useful to examine the landscape in more detail. Doing so allows us to place advanced technical analysis in context.

Some of the earliest research into technical analysis is centred around filter rules (Alexander, 1961; Fama and Blume, 1966; Logue et al., 1978; Sweeney, 1986).¹³ A filter rule simply requires a set percentage move in the daily price of a security from a previous signal price to trigger a buy/sell. For example, a filter may be set to an arbitrary value such as three per cent, and when prices move beyond this in an upward direction, from a preceding low, a buy is recorded. Similarly, when prices decrease by this amount, from a preceding high, a sell is recorded. Filter rules therefore present a seemingly attractive proposition for filtering noisy price data to leave the most economically important price movements.

However, we should remain reticent about these earlier studies which—as might be expected from the date of publication—do not take proper account of risk. Even comparatively recent studies such as Bird (1985) are lacking in this respect. However, an exception to this is provided by Levich and Thomas (1993), who account for risk and use a bootstrap methodology. Profits are found for simple technical trading rules (moving averages and filter rules) employed on five currency pairs. In aggregate, results for this simple strategy are mixed; for example, Sweeney (1986) employs the Sharpe-Lintner CAPM and determines that while returns may not be large, they are not fully explained by risk. However, Corrado and Lee (1992) include a measure of transactions costs and find that while annual returns may be 9.72% greater than a simple buy and hold strategy, these gains are obviated if transactions costs exceed 1.1%.¹⁴

Simple technical analysis strategies also include relative strength, in which buy (sell) trades are entered into for securities, performing strongly (weakly) in prior periods. Levy (1967) initiates discussion in this area, documenting a 32% return

¹³For the UK, Dryden (1970), Cunningham (1973) and Sauer and Chen (1996) evaluate filter rules and find little support.

¹⁴This leads to the conclusion that filter rules are only useful to floor traders who have transactions costs in the region of 0.4%.

gap between NYSE stocks performing the best/worst in prior periods. However, Levy concedes that the superior performance does not refute the random walk due to the lack of suitable methods (at the time) for assessing risk. Jensen and Benington (1970) utilise alphas and betas in replicating Levy's study. With a broad 40-year sample it is found that while the returns from a relative strength strategy are apparently good, they do not exceed the benchmark of a buy-and-hold strategy once risk is accounted for.

Momentum can be considered to be in the same broad category of producing buy/sell signals on the basis of the trend of past price movements. Whilst momentum has only been sporadically connected with technical analysis, much of chart analysis is concerned with looking at trends. In recent years, momentum has become a topic of huge interest in the finance literature. Most significantly, Jegadeesh and Titman (1993, 2001) document profits from momentum strategies over one to three months using US data.¹⁵ The 'momentum effect' also seems to be present in international markets (for example, Rouwenhorst, 1998; Chan et al., 2000).¹⁶

Momentum strategies should be looked at in conjunction with the literature investigating price/return reversal strategies. Where momentum looks to past strength and weakness to continue into the future, reversal (contrarian) strategies anticipate the reversal of current trends. As with momentum, this has become a popular area of research, largely initiated by De Bondt and Thaler (1985). Whilst they focussed on longer-term price reversals, other work has revealed a reversal effect at shorter time horizons, including weekly returns (Lehmann, 1990; Lo and MacKinlay, 1990; Jegadeesh and Titman, 1995). One of the key differentiators between studies showing profits from momentum versus reversal is the time horizon. Chapter 4 looks more closely at this important issue and conducts a detailed analysis of relative strength and momentum effects using a large sample

¹⁵However, one of the earliest studies was presented by James (1968).

¹⁶A more comprehensive examination of the momentum literature can be found in Chapter 4.

of high-frequency data. This is designed to bridge the gap in the current literature, which does not evaluate and investigate these effects at the shortest time horizons.

Moving average based strategies depend upon averaging a 'moving' period of prices prior to the present.¹⁷ By dropping the oldest observation of price in the series with each new observation, this smoothes the price series and buy (sell) trades may be entered into when the moving average turns up (down), or where the faster moving price series crosses the moving average line. More complex strategies may look at two different moving average periods and enter trades where they cross (these often being referred to in the technicians' lexicon as 'golden' and 'dead' crosses).

James (1968) studies monthly moving averages for securities in the period 1926-60. Both unweighted and exponentially weighted moving averages are calculated. The results determine that it is hard to discern out-performance from this trading strategy when set against a simple buy-and-hold alternative. However, only looking at monthly price data can be regarded as restrictive. Investigation into the moving average technical trading strategy has continued in more recent times. Using daily foreign exchange data, Sweeny and Surajaras (1989) compare the trading systems with single and double moving averages as well as filter rules. They find that (on a risk and transactions cost adjusted basis) that the single moving average performs best. Evaluating moving averages up to 200 days in length, Silber (1994) investigates performance in 12 futures markets, including foreign exchange and commodities. For all but three contracts, profits were positive and significant, and the results take account of transactions costs.

More recently, Lee and Mathur (1996a, b) investigate moving average rules applied to daily currency spot rate data from 1988-92. They find, in general, the strategy did not yield significantly positive returns. Optimising the rules and testing out-of-sample did not improve the results. Maillet and Michel (2000) find that

¹⁷Strategies can become more complex by applying a different (heavier) weighting to recent prices, in what is termed the exponential or exponentially weighted moving average.

moving average profits for currency cross rates for a more recent sample appeared to be significant, outperforming a buy-and-hold strategy. The conflicting results of these studies suggests that the success of the moving average is conditional on its specification (weighted versus unweighted, as well as the weighting method) and the sample period being investigated. Research into moving averages continues. More recently, for example, Martin (2001) finds that moving averages constructed up to 30 days for twelve developing country currencies generated significantly positive returns, although this result does not persist after accounting for risk. Olson (2004) finds using daily foreign exchange rate data from 1971-2000 that moving average profits seem to have eroded over time. This is suggestive that the active use of technical analysis may have caused potential gains to be arbitrated away.

Brock et al. (1992) look at a long sample of the Dow Jones Index (1897-1986) and document the success of a simple moving average rule in several sub-periods. These results are merited because Brock et al. account for risk and employ bootstrapping. Even so, they solely look to the Dow Jones Index. There are three important updates to this study that may colour opinion. Firstly, Bessembinder and Chan (1998) propose that the positive abnormal returns they generate from a more widespread study of US indices are due to measurement errors introduced by non-synchronous trading. Secondly, Ready (2002) updates the Brock et al. (1992) study for the period 1987-2000, with the discovery of poor performance of the rule in these years. Further, through a methodology seeking to analyse the *ex-ante* position of traders, they come to the conclusion that the earlier results of Brock et al. is a consequence of data snooping. Finally, Sullivan et al. (1999) determine that while the original Brock et al. results are robust to data snooping, these results cannot be duplicated in an out-of-sample period of 1987-1996.¹⁸

Criticisms of Brock et al. (1992) notwithstanding, there has been continued re-

¹⁸Hudson et al. (1996) provide similar scepticism when repeating the Brock et al. methodology in the UK.

search into moving averages. For example, examining the adaptive moving average which has some degree of response to market volatility (Ellis and Parbery, 2005). Also, the simple moving average is still alive and well, albeit used to investigate new fields such as the profitability in terms of internet stocks (Fong and Yong, 2005). However, neither of these recent papers show sustained profitability from any variants of the moving average.

Looking at studies of international markets adds little to achieving a consensus. For example, Lee et al. (2001) find economically significant returns in four out of thirteen Latin American currency futures investigated. This is hardly overwhelming evidence. In addition, Ratner and Leal (1999) investigate variable length moving averages for emerging markets, finding that three out of ten cases provide significant returns.

While it would be incorrect to identify those mentioned above as the only basic technical analysis strategies, the majority of others, such as oscillators, seek to apply these core indicators in a different fashion. Given that there is a distinctly mixed picture as regards to the overall success of these basic strategies, it is desirable to look towards more advanced strategies and introduce the head and shoulders pattern.

2.3.4 Advanced technical analysis implementations

The broad testing of the above strategies in the literature is largely because of their ease of expression algebraically, and the consequent relatively simple methodology. It is only recently that more advanced patterns have begun to be evaluated by financial economists despite a long history of use by market participants. These complex strategies tend to be “visually orientated”, meaning they involve the recognition of an apparent pattern in past prices. Such patterns are often relatively easy to spot when scanning ones’ eye over a chart of past price histories. However, it is significantly more difficult to program or train (in the case of genetic algorithms or neural

networks) a computer to do the same. This has led to the development of complex algorithms to aid in their detection. A range of methodological approaches have been adopted, including artificial neural networks which can be 'trained' to recognise patterns and smoothing estimators employed in kernel regression. However, this strand of the literature is relatively new and research potential is far from exhausted. This section seeks to present and evaluate the most recent research relevant to this study. The head and shoulders pattern is used in this work as the best exemplar of advanced technical analysis as described above. Furthermore, its history of use by practitioners has been long documented. Accordingly, academic study of the head and shoulders has significant value outside the scholarly community, and beyond simply being a test of market efficiency. However, it is important to grasp what is actually meant by a head and shoulders pattern, and this is addressed below.

While predominated by visually oriented patterns, there are other more advanced technical trading strategies that merit attention. In a similar fashion to patterns, these have only recently started to be evaluated in the literature. An example is point and figure charting, where price moves are represented by 'X's for increasing price moves and 'O's for price declines. While a tool with a long pedigree in the practitioner literature, a tiny body of literature exists. In fact, it appears only two English papers exist on this subject (Elliott and Hinz, 2002; Anderson and Faff, 2008). This existing research is subject to a number of shortcomings. These will be fully examined in Chapter 5, which carries out an investigation into the highly interesting area of point and figure charting using high-frequency price data.

In a similar vein, candlestick charts attempt to allow a full information set to be used in decision making by drawing 'candles' based upon the open, close, high and low prices for a trading day or intraday period. Various formations have been identified by technical analysts. While these are loose patterns due to

their fixed nature over a number of sequential bars, they are relatively trivial to implement programmatically. Again, it is only relatively recently that candlesticks have appeared in the literature. For example, Marshall et al. (2006) construct a candlestick trading strategy for Dow Jones Industrial Average stocks from 1992 to 2002, finding that there is no evidence of profitability. However, whilst some sensitivity analysis is conducted, the main holding period considered is ten days, which may be too short to capture the influence of candlestick trading signals. Marshall et al. (2008a) also assert that candlestick charting is not profitable in the Japanese market over the period 1975 to 2004.¹⁹ Research, however, needs to be extended in this area to allow a full understanding of whether or not a candlestick trading strategy proves profitable over varied holding periods and with intraday data.

This summary of more basic technical trading strategies shows that there is a considerable body of research in some areas, for instance, moving averages. However, a considerable number of technical trading techniques remain under-investigated. Even in the case of moving averages, research continues and is being extended into the myriad different types of averaging techniques that exist. Whilst some of these strategies are more advanced than, for example, basic moving averages and filter rules, and their recent study and evaluation is interesting, they are somewhat different to chart patterns such as the head and shoulders. In particular, these forms of technical analysis depend less on 'visual' identification. Research into pattern recognition has been initiated more recently.

2.3.5 Pattern recognition and the head and shoulders

The head and shoulders falls into the category of price patterns; patterns should be viewed separately to various technical indicators such as moving averages, oscillators and relative strength. Such indicators are based on formulae ranging

¹⁹As candlestick charting is shown to be unprofitable in both Marshall et al. (2006) and Marshall et al. (2008a), neither study proceeds to look at the impact of transactions costs.

from the simple to the relatively complex and provide signals to undertake a trade as their result. In contrast, patterns are visually recognisable formations in a price chart which fit certain geometric characteristics. It is their existence in itself that leads to a conclusion that a buy or sell trade should be entered into.

Clearly, for pattern recognition to be a technical trading strategy, patterns should be repeated over time. It may be that changes in expectations among investors are not immediately assimilated, and patterns form as beliefs become incorporated into prices. The practitioner literature does pause to consider this point, for example:

“A basic principle of technical analysis is that security prices move in trends. We also know that trends do not last forever. They eventually change direction, and when they do, they rarely do so on a dime. Instead, prices typically decelerate, pause, and then reverse. These phases occur as investors form new expectations, and by doing so they shift the security’s supply/demand lines.

The changing of expectations often causes price patterns to emerge. Although no two markets are identical, their price patterns are often very similar. Predictable price behaviour often follows these price patterns.”

(Achelis, 2001, p.245-6)

The head and shoulders pattern stands out as the most important candidate for study for several reasons. First, it has a long history of use, thus going some way towards ameliorating claims of data mining. Second, it is perceived to be a very ‘reliable’ indicator of an impending price movement which can be profitably exploited. For example, Achelis (2001, p.246) describes it as “the most reliable and well-known chart pattern”; Bulkowski (2005, p.406) states that its popularity arises from “its reliability, performance and easy identification.”²⁰ As one of the most

²⁰While this ‘easy identification’ may be possible with the naked eye, as will be seen below, this still presents a serious problem when constructing an algorithm to programmatically find patterns and imitate traders artificially.




Figure 2.1: The head and shoulders top pattern

prominent technical analysis texts for practitioners, Murphy (1999, p.103) describes the head and shoulders as “probably the best known and most reliable of all major reversal patterns”.

As one might guess from the name, the pattern is composed of a central ‘head’ with a ‘shoulder’ on either side. Essentially, a head and shoulders pattern arises out of a series of three peaks—local maximums—in price data, of which the central peak is perceptibly higher than the first or last in the sequence. The basic features of the head and shoulders pattern are illustrated in Figure 2.1. This price formation is more accurately referred to as a head and shoulders *top* and, according to technical analysts, implies a downward price reversal following a prior uptrend. Conversely, technical analysts also recognise the head and shoulders *bottom*, also referred to as an inverse head and shoulders pattern. Shown in Figure 2.2, this is a mirror image of the head and shoulders top. Defined by a series of three troughs (local minima), the central trough is seen to extend further downwards than the ‘shoulders’ on either side.

It should be remembered that we are looking at two patterns: the head and shoulders top and the head and shoulders bottom. As stated above, the occurrence of a head and shoulders top formation is taken to forecast a price decline following a prior uptrend, and the head and shoulders bottom to forecast a price rise following a prior downtrend. The head and shoulders bottom is an exact mirror image of the top formation.




Figure 2.2: The head and shoulders bottom pattern

One of the most important early studies concerned with head and shoulders patterns is produced by Chang and Osler (1995). They detect head and shoulders patterns in foreign exchange data for the spot rates of six currencies against the dollar in the period 1973-1994. It is found that out of the six, only two currencies (the Mark and the Yen) offer significant profits. These profits are robust to the inclusion of transactions costs (reflected in the analysis as a moderate bid-ask spread) and do not appear to be a reward for bearing systematic risk. The latter point is demonstrated by estimating the beta for the excess returns owing to the head and shoulders strategy. However, the conclusion is that if a head and shoulders strategy were employed in all the currencies simultaneously across the sample period then this is sufficient to capture abnormal returns in aggregate.

Chang and Osler recognise some criticism to their own study; in particular, the consideration of only a narrow sample drawn from foreign exchange markets. Furthermore, they are ambiguous about the length of period in which patterns form. In reality, traders may look for these patterns developing in a short space of time using tick-by-tick intraday data or over a longer period of months or even years using weekly data. Traders may be particularly interested in patterns forming over the shortest time horizons in the case of foreign exchange markets, which are dominated by volatility and short term trading. Indeed, Chang and Osler find that under their methodology, positions are taken in a particular currency once or twice a year, on average. Many traders in the foreign exchange market employ a

significant amount of leverage, and would be unlikely to want exposure from a position taken over a long time horizon. They would therefore be unlikely to use these infrequently occurring signals, preferring a greater number of smaller trades. Whilst it is contended that this is in accordance with the practitioner literature, it is more the case for stock rather than foreign exchange. Considering this criticism, they would have perhaps been better to also investigate stock indices and/or individual stocks.

Head and shoulders patterns are defined by a series of localised maxima and minima in price data (peaks and troughs) as shown in Figures 2.1 and 2.2. It is therefore the case that the method used to detect these local extrema conditions the ability to isolate head and shoulders patterns. Chang and Osler use a filter rule type methodology. Peaks (troughs) are defined as a local maximum a set percentage higher (lower) than the preceding maximum (minimum). They 'scan' the data ten times, using different 'cutoff' percentage values each time and merge the results to get an overview of peaks and troughs. This approach is problematic because the patterns that are detected may be unbalanced. For instance, the left shoulder and head may owe to peaks detected with a relatively high cutoff value, yet the right shoulder and completion of the pattern may be due to a peak detected with a relatively low cutoff. This may tend to give patterns that are detected computationally, yet would be unrecognisable to professional technical traders looking for such formations.

Whilst Chang and Osler put in place restrictions to rule out patterns which appear to be too unbalanced, and do not have the correct 'symmetry', this method is still open to question. It is possible to make improvements on using a filter-rule type method to detect peaks and troughs, as will be seen below.

However, Chang and Osler (1999) seek to address some of the criticisms of their earlier paper. They determine that the original results *may* have been inefficient. More importantly, however, they provide evidence that the more complex head

and shoulders strategy is surpassed in terms of returns by simpler strategies. The implication is that advanced technical analysis is not worthwhile.

Lo et al. (2000) provide one of the most comprehensive recent studies on advanced technical analysis. They look at a variety of patterns including head and shoulders, but also less popular and less well known patterns including broadening tops/bottoms, rectangle tops/bottoms and double tops/bottoms. The most interesting development in this research is the use of non-parametric kernel regression to locally smooth price data. In this methodology, a bandwidth is chosen that is designed to eliminate excess noise from the series, whilst retaining economically useful information about the underlying trend. It is possible to then take the derivative of this series and use it to identify turning points which can be translated into localised maxima and minima—the building blocks of advanced technical analysis with chart patterns. This is a far more attractive methodology than the filter-rule type analysis employed by Chang and Osler (1995). Once localised maxima and minima are found, a rule-set can be applied to find patterns. The study further employs bootstrapping to evaluate the significance of results. The empirical work in the present chapter and Chapter 3 makes a significant modification to Lo et al.'s kernel methodology. In doing so, the methodology adopted here is better suited to technical analysis and eliminates an important element of subjectivity in their approach.²¹

The stated aim of Lo et al. (2000) is not to discover the returns accruing from pursuing a technical analysis strategy directly. Instead, they seek to assess the informational content of patterns through a comparison of the distribution of returns conditional upon a particular strategy with unconditional returns. The idea is that “[i]f technical patterns are informative, conditioning on them should alter the empirical distribution of returns; if the information contained in such patterns has already been incorporated into returns, the conditional and unconditional

²¹This concerns the *ad-hoc* modification of optimised bandwidth, which risks torturing the data. This is discussed further below, where an alternative methodology will be presented.

distribution of returns should be close" (p. 1726). The net result is that it is not known whether using price patterns constitutes a profitable trading strategy.²² This shortcoming is comprehensively addressed by the empirical work in this study which, in the context of a trading strategy, presents returns from holding periods of 1 to 60 days.

While using individual NYSE/ AMEX and NASDAQ stocks from a long time period from 1962 to 1996, a major shortcoming of the Lo et al. (2000) study is that it only uses 10 (randomly selected) stocks from each market capitalisation quintile in each period. Thus, sampling with replacement, there are only 50 stocks for each sub-period. However, perhaps the greatest weakness is that, in spite of the stated aim to assess informational content of patterns, only the one-day continuously compounded return is examined at a set number of days after patterns have formed. Traders are unlikely to exit the trades recorded by Lo et al.'s study after just one trading day. In their study, patterns can form over 35 trading days. The practitioner literature tells us that traders will often hold a trade for a time similar to the formation period of the head and shoulders pattern. Thus, looking solely at 1-day returns is unrealistic and a severe shortcoming of this study. Given these problems, it is surprising that Lo et al. find that a range of their patterns, including the head and shoulders, possess informational content, and that these findings are relatively robust.

Dawson and Steeley (2003) replicate Lo et al.'s methodology extremely closely for the UK. They draw their sample as a sub-set of the constituents of the FTSE100 and FTSE250 indices from 1986 to 2001. However, just 15 stocks from each size quintile are selected in each sample period, for a total of 75 over the study. Results are similar to those of Lo et al. (2000) - that the head and shoulders pattern appears to have some predictive power but this is not quantified. The sample is both shorter and less broad than that used here. In this study, we use the largest 350 stocks by

²²As the stated aim is not to examine profitability, transactions costs are not accounted for.

market capitalisation over the period 1980-2003. Most importantly, Dawson and Steeley's study is subject to the same criticisms made above concerning Lo et al. (2000). Therefore, whilst the results from this study are interesting, they still give no clear picture of the profitability of the head and shoulders in the context of a trading strategy. Furthermore, the patterns identified are unlikely to accord with those recognised by traders. Accordingly, the empirical work presented here represents a significant contribution, although we both work with (different periods of) UK data.

Very recently, Savin et al. (2007) provided a further investigation of head and shoulders patterns using US data. In the above discussion, two of the disadvantages of Lo et al. (2000) were the lack of ability to assess the profitability of a head and shoulders trading strategy (as they only look at 1-day returns) and the relatively small sample. Savin et al. (2007) make some advances in attempting to address these issues in looking at returns over 20, 40 and 60 days together with using a sample based on the S&P500 and Russell 2000. They find little support for the profitability of a head and shoulders based trading strategy, although the significance of excess returns suggests that the patterns do have predictive power. In the cases where profitability is evident, this is often subsumed by transactions costs. For instance, the excess return to the S&P 500 trading strategy, net of transactions costs, remains positive only at 60 days.²³ The study also seeks to look at risk-adjusted returns using the conventional three-factor model and a four-factor model which also includes a momentum factor. Results suggest that, in some cases, the head and shoulders pattern is profitable after risk and transactions cost adjustments. There is, however, no conclusive evidence of profitability overall as part of a stand-alone trading strategy.

However, whilst this study is valuable, there are some crucial gaps. The most noteworthy is that the study only looks at head and shoulders tops and not bottoms.

²³To do this, Savin et al. (2007) look at the one-way break-even cost, comparing to one-way transactions costs identified in previous research.

Technical analysts view the head and shoulders as a symmetrical pattern that has predictive power that can be harnessed for both long and short trades. Further, Savin et al. (2007) adopt the kernel regression method of Lo et al. (2000), but examine different multiples of the bandwidth. Given the discussion above, that arbitrarily altering an optimised bandwidth has the potential for data mining, this is not an ideal approach. The empirical work in this chapter employs a newer method for using kernel regression with a local optimised bandwidth. Further, this study introduces the trade lag and looks at both head and shoulders tops and bottoms.

Lucke (2003) looks at head and shoulders patterns in the spot rate of five currencies relative to the US dollar. The sample lengths vary for the various currencies, but broadly reflect an approximate 20-25 year period for each. This study is particularly interesting as it adopts some of the methodology of Chang and Osler (1995) and looks at several different implementations of pattern geometry. The results are uninspiring from a technical analysts' viewpoint: it is found that "[r]eturns to SHS trading are not significantly positive - and if there is any evidence for non-zero returns at all, then it is evidence for negative returns" (Lucke, 2003, p.39). Given the absence of profitability, transactions costs are not considered. However, there is also no treatment of risk, which may help to explain the significant negative returns that are found.

However, studying daily spot rates may be an inappropriate use of the head and shoulders patterns. In particular, if we postulate that the head and shoulders pattern develops out of market inefficiency in compounding investors' expectations, then the highly liquid forex market is perhaps the harshest proving ground. Given this, it may also be that tick-by-tick data would be more appropriate for this study. Furthermore, Lucke uses business cycle turning-point detection methods to find peaks and troughs in data, owing to Bry and Boschan (1971). It would be more appropriate to adopt a more modern methodology of kernel smoothing, as used

in this study.

2.3.6 Technical analysis in the markets

It is important to establish that traders do, in fact, use technical analysis to inform their decisions. A number of surveys of traders have been exhibited in the literature. Clay (1925) indicates just how early technical analysis was in popular use. From *Moody's Investors' Service*, Clay described several methods of forecasting stock prices but determined that “the most popular method of forecasting is chartreading” (p.245). Despite the author being sceptical of its benefits, it is further testament to the widespread use of technical analysis over a long period.

The use of technical analysis in US commodity futures was first recorded in the academic literature by Smidt (1965). Even before this, Stewart (1949) records the use of strategies akin to technical analysis in Chicago futures trading.

The Group of Thirty (1985) conducted a wide-ranging early study concerning the functioning of foreign exchange markets. Spread over 12 countries, 40 banks and 15 securities houses were queried. Technical analysis appeared to be almost universally popular, with 97 per cent of banks and 87 per cent of securities houses reporting its use.

Frankel and Froot (1990) focus on foreign exchange forecasting services. They provide results from *Euromoney* magazine from services surveyed between 1978 and 1988. At the start of the sample period, in 1978, out of a total of 23 firms surveyed, only three report that they used technical analysis. This is in comparison to 19 using fundamental models. By 1988—where 31 firms are queried—18 report usage of technical analysis, 7 report the use of fundamentals and 6 state the use of both technical and fundamental analysis. According to the data, the swing in favour of technical analysis seems to have started in 1983. This may be because of increased computational availability and reduced transactions costs making more frequent trading viable. While the sample is relatively small, it is a useful study in

showing the changing use of technical analysis.

In addressing the question of the use of technical analysis amongst traders, one of the most useful and interesting surveys is provided by Taylor and Allen (1992). The authors composed a questionnaire that was dispatched to foreign exchange dealers in London with this issue in mind. The survey was designed to elicit responses as to both how technical analysis was employed and how dealers regarded its usefulness.²⁴

One interesting result is that out of the sample of 213 responses, two fifths employed in-house economists, and out of these 38.5 per cent took positions (as opposed to their role being purely advisory). While only a quarter per cent employed in-house technical analysts, 45.5 per cent of these took positions in foreign exchange. Clearly some organisations had a definite preference for technical analysts' forecasts over fundamental style forecasts produced by economists.

Usefully, Taylor and Allen break down the influence of technical forecasts on dealers' activities by time horizon. One might reasonably expect that as technical analysis concerns itself solely with past price history and other summary statistics that it would be most useful at short time horizons. This is because we would expect in even a moderately efficient market that any abnormal profits would be speedily arbitrated away. The study confirms this with 90 per cent of respondents using some information from technical analysis from intraday to one week horizons. 60 per cent regarded technical information to be at least as informative as fundamental information. The survey also notes that over all time horizons there were some respondents who never employed fundamental analysis and solely made trading decisions based upon technical forecasts.

In a more recent survey of London foreign exchange dealers, Cheung et al. (2004) find similar support for technical analysis as a valued tool. Menkhoff (1997) shows that technical analysis is also heavily used by foreign exchange dealers

²⁴The survey data was collected in 1988.

in Germany. Respondents to a survey indicated that technical analysis had an effect on trading decisions at timespans from intraday to 2-6 months. One finding, which is surprising given that technical analysis has a long history, is that younger respondents tended to have a stronger preference for technical analysis over their older colleagues.

Conducted in 1995, Lui and Mole (1998) undertook a questionnaire survey of Hong Kong foreign exchange dealers. They ascertain yet again that technical analysis is regarded as important, particularly at shorter time spans of up to six months.

Cheung and Wong (2000) also report the results of a 1995 study seeking to investigate market microstructure issues. From the survey of individual traders working on the Hong Kong, Tokyo and Singapore exchanges, it is found that technical analysis is important. Indeed, “[a]bout 40% of respondents say that technical trading is the major factor in determining exchange rates in the medium run.” (p.411) Perhaps most surprisingly, from an efficient markets point of view, is that “Even in the long run, 17% of traders still believe technical trading is a significant determining factor.” Cheung and Wong define the long run as beyond six months.

Further evidence of the use of technical analysis in futures markets is supplied by Brorsen and Irwin (1987) in terms of public futures funds’ advisory groups. Cheung and Chinn (2001) survey foreign exchange traders in the US and also find support for technical analysis.²⁵ Indeed, 30 per cent of respondents classified themselves as primarily trading using technical analysis signals. A similar result was found when Cheung et al. (2000) surveyed foreign exchange traders in the UK, with 33 per cent of respondents identified as technical traders. Both of these surveys indicated an increase in the use of technical analysis compared to five years previously.

²⁵This survey was conducted in 1998.

Gehrig and Menkhoff (2006) survey 200 foreign exchange dealers and international fund managers in Germany and Austria.²⁶ They include questions designed to invite a response as to the preference of dealers for fundamental factors, order-flow and technical indicators. The results show that “technical analysis dominates foreign exchange and most FX traders seem to be chartists now, but [m]ost professionals use charts and fundamentals in a complementary manner” (p.3). Recipients concurred with those from other surveys (above) in reporting that charts are dominant primarily at short time horizons.

Oberlechner (2001) provides results of a survey of how not only foreign exchange traders but also financial journalists regard technical and fundamental analysis.²⁷ The questionnaires and interviews conducted in Frankfurt, London, Vienna and Zurich again show that technical analysis is widely used, in particular at short forecasting horizons. Its use seems to have increased since Taylor and Allen’s 1992 survey, where data was collected in 1988. Interestingly, traders place more emphasis on technical analysis as a viable forecasting tool than do journalists. In general, it was shown that traders in Vienna and Zurich used technical methods more than their counterparts in London and Frankfurt.

Research on the use of technical analysis in making decisions about equity trades is scarcer, although Arnsward (2001) also finds evidence of technical analysis use in making short-term investment decisions of up to eight weeks among German mutual funds. A large number of responses allowed methods of analysis and forecasting to be ranked, with technical analysis achieving second place, behind fundamental analysis and ahead of portfolio optimisation and econometric models.

Similarly focussing solely on collective investment vehicles, Menkhoff and Schmidt (2005) survey mutual fund managers in Germany.²⁸ While not specifically looking at technical analysis, there are some useful insights. Most fund managers

²⁶Questionnaires were sent out in 2001. Usefully, the results can be compared to a similarly designed survey in 1992 to assess the changing importance of technical analysis over time.

²⁷Data was collected in 1996.

²⁸Questionnaires were sent out in 2002 to 64 fund management companies.

used contrarian, buy-and-hold and momentum strategies. When participants were asked about what information sources they used, conditional on the type of strategy being pursued, it was found that those pursuing momentum strategies utilised technical indicators the most. A borderline result was achieved for contrarian strategies.

It therefore seems that technical analysis has been a feature of foreign exchange trading and among banks and mutual funds for some time. The evidence presented above suggests that its influence has increased and not weakened. There is less survey data available for the use of technical analysis in equity trading. However, much of the practitioner literature is focussed on technical analysis in stock trading. Even so, the survey data shows that technical analysis is phenomenally popular amongst traders, in particular at short time horizons. Some data even suggests that its importance is increasing. This supports the interesting nature of technical analysis in the context of academic study.

2.3.7 Literature review conclusions

The above sections have sought to evaluate the opinions that academic literature takes on technical analysis strategies. First, it is useful to draw some general conclusions. Most importantly, the literature is characterised by fragmentation with little commonality in methodology or approach. It has been shown that the early literature largely focuses on basic technical analysis strategies and, once risk is taken into account in slightly later studies, it is not supportive of the generation of economic profits. More recently, there has been something of a renaissance in interest in technical analysis, accompanied by the application of more advanced statistical techniques and methods, and the investigation of complex patterns such as the head and shoulders.

Having highlighted the fragmented nature of the literature and the large number of gaps in markets examined, it is interesting to take the body of research

further with an investigation into the head and shoulders pattern in the context of the UK. The following section describes the methodology that is used to achieve this.

2.4 Data and methodology

2.4.1 Nature and breadth of data

It is highly desirable to use a large dataset based around individual securities for two main reasons. First, a small number of stocks gives less concrete grounds for inference on the value of the head and shoulders. Second, much of the practitioner literature relates to head and shoulders patterns occurring in price charts of individual stocks, rather than indices. The review of the literature above revealed that the comparatively few detailed studies of advanced technical analysis have often used narrow data sets. To reiterate, Brock et al. (1992) look at 90 years of data, but only for the Dow Jones Industrial average. The landmark study of head and shoulders patterns from Chang and Osler (1995) focuses on the spot rate for a small sample of six currencies against the dollar. Only more recently, with less costly and more powerful computational capabilities, have larger data sets been used. For example, Lo et al. (2000) look at just over thirty years of individual US stock data from NYSE/AMEX and the Nasdaq. However, while not squandered, this sample is not put to its full use. Presumably in an effort to reduce computational time and data collection and management, Lo et al. randomly draw stocks from this larger sample to end up with just fifty stocks per period under investigation. Dawson and Steeley (2003) adopt a similar approach for the UK, but only select a random sample of 75 stocks per five year period under investigation from FTSE 100 and 250 constituents. In addition, their methodology and approach to the chart patterns is considerably different to that employed here.

In contrast to previous work, a large dataset of individual stock prices will be utilised for the United Kingdom. The Datastream research service was used for the collection of price and market capitalisation data.²⁹ The sample period is January 1 1980 to December 31 2003, representing 24 years of daily data, comprising 4983

²⁹Thomson-Financial (2005).

stocks. The database contains 31,205,920 observations in total. Crucially, dead stocks are included to avoid any survivorship bias.³⁰

A portfolio of the largest 350 stocks by market capitalisation is generated annually, from all the stocks traded on the London Stock Exchange at that point in time. All stocks for which there is a price available on the formation date are considered. The largest 350 stocks can be considered a close proxy for constituents of the FTSE 350 (the FTSE 100 and 250 together).³¹ The study does not suffer from survivorship bias as we include all stocks in the data set. Therefore, if a stock ceases to trade then returns will be calculated appropriately, with a 100% loss for long trades and a 100% profit for short trades.

2.4.2 Identifying peaks and troughs in price data

The name of the head and shoulders pattern comes from its nature as a sequence of peaks and troughs which visually appear to approximate the head and shoulders of a human being. While a pattern could conceivably be seen over just six trading days, as shown in Pattern 1 in Figure 2.3, this is atypical. As identified above, traders may perceive a head and shoulders pattern forming over a long time period, whether this be on a weekly, daily or intraday price chart. Consider the example of Pattern 2 in Figure 2.3. Pattern 1 forms from t_1 to t_7 (6 trading days), but Pattern 2 is between t_8 and t_{17} (9 trading days). For any time longer than six days, a set of extrema (local maxima and minima) serve to define the pattern. In order to identify head and shoulders patterns a suitable method of detecting these peaks and troughs in the data is needed.

The apparent ease with which this can visually be done *ex post* is deceptive in two ways. First, it is problematic to use a computer to find peaks and troughs

³⁰Considering the large sample and the need to retrieve random access data, in order to compute returns based on exit criteria varying by trade, for use in the detection of head and shoulders patterns, it was necessary to develop a database system.

³¹This is an approximation because some stocks may have been excluded from the index due to free float requirements or other considerations by the selection committee at FTSE International.

without a rigorous definition of what constitutes a local maximum/minimum in the data. Second, it is not clear when a peak or a trough will occur until it has fully formed. However, if we are seeking to replicate traders' activity we cannot introduce the look ahead bias that an *ex post* filtering of price data would involve. In practical terms, this latter problem is somewhat ameliorated by the specification of a head and shoulders pattern itself; a trade can be entered before the last trough (peak) for a head and shoulders (inverse head and shoulders) pattern has formed. In terms of Figure 2.3, a pattern can be classified as complete at t_6 for Pattern 1 and t_{16} for Pattern 2. This is the approach adopted by Lo et al. (2000).




Figure 2.3: Formation of head and shoulders patterns

The method selected to identify peaks and troughs in the data is by the use of a smoothing estimator and kernel regression. The essence of visual price patterns is the extraction of a non-linear formation from noisy price data. Lo et al. (2000) point out that smoothing estimators are well suited to this task as they average out the noise and allow us to extract useful signals.

2.4.3 Kernel smoothing

If the detection of peaks and troughs in price data is a prerequisite for identifying price patterns then it is necessary to identify a suitable method for isolating these local maxima and minima. There are a number of possible ways of achieving

this. For instance, Bry and Boschan (1971) developed a method of programatically detecting turning points in time-series data, aimed at isolating business cycles. Their methodology used a 'moving window' which essentially looks at identifying a local maximum (minimum) depending on whether an individual data point is higher (lower) than a pre-specified number of surrounding points. Results from a series smoothed using a moving average are aggregated with an unsmoothed series to isolate turning points. However, this methodology is somewhat subjective as it requires the widow width, length of moving average and criteria for removing turning points which occur in quick succession to be pre-specified.

A further possibility is to compute a moving average for a price series and then look for times at which this moving average changes direction. However, there is still the need to specify the length of the moving average. Furthermore, it would seem unlikely that a single fixed-length moving average could adequately capture times of large price changes clustered together where many peaks and troughs occur in close proximity.

Employing kernel regression smoothing addresses such disadvantages. Kernel regression is a nonparametric method which allows us to fit a curve to non-normal, noisy data. It defines itself as a nonparametric methodology by the use of a kernel to decide the weight that is placed upon each data point for calculation of the smoothed value.³² By smoothing out a series, we hope to be able to pick up non-linear relations by a complex averaging procedure. Kernel regression is only one way of achieving this; spline functions, wavelets and nearest-neighbour estimators are some of the alternatives.

Smoothing estimators have become popular in the finance literature; for example, their use by Diebold and Nason (1990) and Meese (1990) in investigating nonlinearities in the foreign exchange market. Both apply nearest-neighbour tech-

³²Note that we are not interested in the nonparametric aspects of kernel regression *per se*, but rather in its application to smoothing. It is the detection and evaluation of chart patterns that are the main subject of this study.

niques, although neither has great success in application. Gencay (1999) shows more positive results, looking at the linear and non-linearity predictability of the spot exchange rates arising out of a moving average indicator. The result is that the non-linear forecasts from nearest-neighbour and feedforward regressions dominate those from random walk and GARCH(1,1) models.

From the viewpoint of technical analysis, Lo et al. (2000) presented the case for employing smoothing methods to aid in pattern recognition. Lo et al. justify the relevance of smoothing estimators in the study of technical analysis on the basis that they “extract non-linear relations $\hat{m}(\cdot)$ by “averaging out” the noise. Therefore we propose using these technical estimators to mimic and, in some cases, sharpen the skills of a trained technical analyst in identifying certain patterns in historical price series” (p.1708). One particularly attractive feature of smoothing estimators is that they approximate the way that humans visually extract patterns from noisy data (Poggio and Beymer, 1996).

Whilst this research accepts Lo et al.’s thesis that kernel regression is a useful way of extracting non-linear patterns from noisy data, we look at two alternative kernel-based smoothing approaches. Whilst similar, this study contends that these methods are superior to the approach adopted by Lo et al, who use the Nadaraya-Watson estimator (Nadaraya, 1964; Watson, 1964).³³ The methodology and reasoning are discussed below.

Beginning with a series of security prices, P_t , we can specify a model of the form

$$P_t = m(X_t) + \epsilon_t \quad i = 1, \dots, n \quad (2.1)$$

where $m(\cdot)$ is the unknown mean regression function and ϵ is a random variable (white noise). In this case, $m(X_i)$ is a function of time, effectively the kernel

³³A good exposition of the Nadaraya-Watson estimator, and kernel smoothing in general, can be found in Hardle (1990) and Eubank (1999).

smoothed price series.

A central desire of employing nonparametric regression is to smooth the price series so that maxima and minima can be isolated, but while taking account of the noisy nature of price data. To achieve this, a weighting scheme is needed such that prices near a particular point in time receive larger weights. The approach employed by Lo et al. (2000) is to use the Nadaraya-Watson kernel regression estimator. This is defined as

$$\hat{m}_{NW}(x) = \frac{\sum_{t=1}^T P_j K_h(x - X_t)}{\sum_{t=1}^T K_h(x - X_t)} \quad (2.2)$$

where $K(\cdot)$ is the kernel (see below) and h is a positive parameter. h is often referred to as the bandwidth, and it is this value that determines the smoothness of the resulting estimate. This is because if a large value of h is chosen the average is computed over a large number of points around X_t . Alternatively, if a small value of h is chosen then only the closest points around X_t are averaged. Thus, by reducing the value of h , the fitted curve more closely follows the original price series. On the other hand, a large h results in a curve that does not follow the original series as closely, yet may better illustrate the important local trends in prices. The term trend is used here in the context of technical analysis. It is the change between localised uptrends and downtrends that gives rise to the key points of the head and shoulders pattern.

To reinforce and illustrate the importance of the bandwidth, examples of different bandwidths applied to a stock price series are presented and discussed below. Before this, however, a little more needs to be said concerning the kernel function itself, $K(\cdot)$. If the bandwidth, h , can be thought of as dictating the 'size' of the weights then the kernel, K , dictates the 'shape' of how these weights are applied around an individual observation. The most frequent choice of kernel in empirical studies is the Gaussian kernel, which is

$$K_h(x) = \frac{1}{h\sqrt{2\pi}} \exp\left(-\frac{x^2}{2h^2}\right) \quad (2.3)$$

There are many choices of kernel, some of the more popular selections being the Epanechnikov and Gaussian kernels. Figure 2.4 shows smoothing of stock price data using four different kernels (the Epanechnikov, the Gaussian, the Triangular and the Rectangular.) Looking at the smoothed series against the original British Airways (BAY) share price, which is used as the exemplar, it can be seen that each kernel provides a slightly different representation of the data. There is evidence that the choice of kernel makes little difference to the results of kernel smoothing. For example, Silverman (1986, p.43) notes that, on the bases of the integrated mean square error, that “there is very little to choose between the various kernels”. Further analysis of this issue can be found in Hardle (1990). In consequence, the Gaussian kernel is adopted here.

The far more important choice is in selecting the bandwidth, h . Should this be too large, then useful information is lost by ‘over-smoothing’. If it is too small, then too much noise remains from the original series. To illustrate this, Figure 2.5 shows the price of British Airways against the series smoothed with $h = 1$. Contrast this with Figure 2.6 and Figure 2.7 for $h = 2$ and $h = 5$, respectively. It can be seen that as the bandwidth increases, the smoothed series follows the original less closely. A method of selecting an optimal bandwidth, such that over-smoothing and under-smoothing are avoided, is needed.

A common approach is to use cross-validation. This is sometimes referred to as the ‘leave-one-out’ method, as the first step is to omit one observation and estimate $m(\cdot)$ at x_j , i.e.

$$\hat{m}(x_{h,j}) = \frac{1}{T} \sum_{t \neq j} w_t(x_j) P_t \quad (2.4)$$

After obtaining this smoother for $j = 1$, the same smoother is calculated for




Figure 2.4: Examples of the application of various kernels.

$j = 1$ to $j = T$. The bandwidth, h , is then chosen so as to minimise the cross-validation function, defined as

$$CV(h) = \frac{1}{T} \sum_{t=1}^T (P_t - \hat{m}_{h,t})^2 \quad (2.5)$$

Lo et al. (2000, p.1714), who adopt this approach, state that “the bandwidths obtained from minimizing the cross-validation function are generally too large for our application to technical analysis.” Therefore, the data is smoothed too much and some important peaks and troughs are not identified; in other words, potentially useful data is discarded. This is a balance as the data should be smoothed sufficiently to allow head and shoulders patterns to form over a long enough time period, yet not prohibit what could clearly be visually identified by traders as a pattern in the local extrema.

Lo et al. (2000) choose to use an ad-hoc solution of using a bandwidth of 30% of the optimised value (i.e. $0.3 \cdot h^*$, where h^* is derived from minimising the cross validation function). However, selecting 30% of the cross-validated bandwidth is unsatisfactory because of its subjectively. The cross-validation function is used to derive an optimal bandwidth for the smoothing process by minimising the mean integrated square error, yet this optimised bandwidth is adjusted in their study. It could be argued that the arbitrary alteration of bandwidth is symptomatic of data mining.

This study seeks to overcome this criticism by using different methods of kernel estimation. Having noted above that the choice of the specific kernel (such as Gaussian or triangular) is not of crucial importance in the smoothing process, the manner in which this kernel is applied is considerably more important. Lo et al. (2000) use the Nadaraya-Watson kernel estimator, described above. An alternative choice, adopted here, is the Gasser-Muller kernel estimator (Gasser and Muller, 1979). Comparative studies of the Nadaraya-Watson and Gasser-Muller estimators reveal that both methods have advantages and disadvantages (Chu and Marron,

1991; Jones et al., 1994). However, a crucial benefit of changing the approach taken by Lo et al. (2000) and selecting the Gasser-Muller kernel estimator is that it is possible to use a more appropriate method for selecting an optimal bandwidth.

Formally, the Gasser-Muller estimator is defined as

$$\hat{m}_{GM}(x) = \sum_{t=1}^T \int_{s_{t-1}}^{s_t} K_h(x-u) du P_t \quad (2.6)$$

where $s_t = (x_{i-1} + x_i)/2$.

Park and Marron (1990) note that when applied to real-world data, the performance of the basic cross-validation approach has “often been disappointing” (p.66). One alternative is to use the so-called plug-in approach, advocated by Park and Marron and also Sheather and Jones (1991), who also demonstrate that using the plug-in method leads to considerably less variability. Both cross-validation and the plug-in approach seek to minimise integrated square error (MISE). However, the plug-in approach starts with an approximation of the MISE and iteratively minimises it. There are 11 iterations in the approach outlined by Gasser et al. (1991), who present an iterative plug-in approach for selecting a global bandwidth, based on the Gasser-Muller kernel estimator. This approach is referred to as ‘global’ because a constant bandwidth is used to smooth over the sample.

While Gasser et al. (1991) provide evidence that this approach is attractive compared with cross-validation, perhaps the main driver for the use of this technique here is that it is faster to compute than cross-validation. When examining a very large amount of rolling windows to detect head and shoulders patterns it was found that this produced a considerable time saving. The use of this approach made it feasible to evaluate a large number of simulated series when performing bootstrapping, in order to shed light on the significance of profits accruing from the head and shoulders.

However, it can still possibly be considered sub-optimal that we are using one bandwidth to smooth all data, particularly given the volatility-clustering often

seen in financial time series. Following Brockmann et al. (1993), Herrmann (1997) develops a method of using a locally varying bandwidth instead of a single globally optimised bandwidth. From the perspective of dealing with stock price data, it is particularly relevant that this approach is better at dealing with heteroscedasticity. Furthermore, Brockmann et al. (1993) state that “the estimator can adapt to the structure of the regression function, smoothing more in flat parts of the curve and less in peaky parts” (p.1302). Given that the head and shoulders pattern is defined by peaks and troughs, this approach has an obvious attraction. Accordingly, this locally optimised approach is the second to be evaluated in this study.

Application of kernel smoothing

Kernel smoothing is important in this chapter in terms of its application to detecting maxima and minima. To illustrate how the approach performs, Figure 2.5 shows a plot of two years of daily data on the share price of British Airways (blue line). The figure also shows the fitted curve that is obtained by smoothing with the Gasser-Muller kernel estimator with a bandwidth set to 1 (i.e. $h = 1$). The box-out in the lower right provides a ‘zoomed-in’ view. By contrast, Figure 2.6 shows the same price series but this time a bandwidth of 2 ($h = 2$) is chosen. It can be seen that the fitted curve follows the original series less closely than with $h = 1$. This is even more the case with $h = 5$, as shown in Figure 2.7.

Figure 2.8 again shows the smoothed series derived from the British Airways share price. However, this time bandwidth is set by using the global optimisation technique described above. In this case, bandwidth is found to be $h = 2.16$. For the purposes of comparison, the bandwidth derived under cross-validation with the Nadaraya-Watson estimator is 11.686. This would be transformed to 3.51 by Lo et al.’s 0.3 adjustment factor. Using the global plug-in approach, the optimised bandwidth is 2.16. This is much closer to the 3.51 bandwidth that Lo et al. would have employed for this particular example. However, this is achieved without




Figure 2.5: Bandwidth selection




Figure 2.6: Bandwidth selection




Figure 2.7: Bandwidth selection

having to make an arbitrary and subjective adjustment.

Figure 2.9 gives an example of smoothing with a locally optimised bandwidth following Herrmann (1997). Whilst the smoothed line is superficially similar to global optimisation it is possible to see the effect of localised optimisation; the box-out aids visual inspection by zooming in on a portion of the chart. In particular, notice the different extrema points identified, which are highlighted in green.

After kernel regression, the smoothed series needs to be translated into peaks and troughs to use in identifying head and shoulders patterns. This can be achieved with the signum (sign) function

$$\operatorname{sgn} x \begin{cases} -1 : x < 0 \\ 0 : x = 0 \\ 1 : x > 0 \end{cases}$$

employed on the derivative of the smoothed series. Therefore, at times when the sign of the signum function changes from +1 to -1, a local maximum has been discovered. Conversely, where the sign changes from -1 to +1, a local minimum can be recorded. Figure 2.10 illustrates this process with h (h is bandwidth) set to 2. The top panel shows the raw and smoothed series for the British Airways price, as in previous figures. The centre panel shows the first derivative of the smoothed function. The bottom panel graphically represents the signum function. As it fluctuates between -1 and +1, it identifies when a peak or trough is recorded.

2.4.4 Detecting head and shoulders patterns

Once peaks and troughs in the data are identified they can be used with the definitions of the head and shoulders patterns, above, to identify points at which a trader would buy and sell. It is important to highlight that special care is taken to avoid any look-ahead bias. Analysis is performed on a rolling basis such that each trading day in the sample is treated separately, as if a trader was coming afresh to




Figure 2.8: Bandwidth selection - globally optimised




Figure 2.9: Bandwidth selection - locally optimised




Figure 2.10: Determining peaks and troughs with the use of the signum function. $h = 2$.

the market on that day. In so doing, the program only works with the information set that the analyst would have available on a particular trading day. The steps that are followed in order to identify patterns can be summarised as:

1. At the start of each calendar year, identify the sample of stocks to be studied as outlined above. These are the 350 largest stocks traded on the London Stock Exchange by market capitalisation.
2. Taking 35 trading days of daily data, identify the peaks and troughs (local maxima and minima) by performing kernel regression with globally and locally optimised bandwidths and employing the sigmoid function to isolate turning points.
3. Using these peaks and troughs, apply an algorithm (see below), and identify where head and shoulders patterns occur and log a buy trade for inverse head and shoulders patterns and a short sale for head and shoulders tops.
4. Record the price on exit, and calculate the continuously compounded return. The three-month gilt is used to obtain excess returns, by subtracting its continually compounded return over the same holding period. Six exit points (1, 5, 10, 20, 30 and 60 days) are evaluated.
5. Calculate the results based on the 'trade lag' filter (see below).

It is important to highlight the rolling window method approach that is used. As above, windows of 35 days of data are taken. These overlapping windows are used for each security in the dataset. Figure 2.11 provides a graphic representation of the operation of rolling windows. In this example, there are 40 observations of daily prices, from t_0 to t_{40} . The first rolling window runs from t_0 to t_{34} , the second from t_1 to t_{35} and so on. If a head and shoulders pattern is found within a window, the trade (sell for a head and shoulders and buy for an inverse head and shoulders)




Figure 2.11: Head and shoulders pattern detection using the rolling windows

is entered at the following day's closing price. For instance, if a pattern is found in the first rolling window, from t_0 to t_{34} , a trade is recorded at t_{35} .

This methodology has two advantages. First, it means that we operate in a fashion similar to technical analysts who view each day of data 'as it comes' and look for the occurrence of new patterns. Second, as noted by Lo et al. (2000) and Savin et al. (2007), this means that we identify all the possible chart patterns that would be available to traders. If windows were sequential instead of overlapping we would only capture patterns that completed in one window. Suppose we had two 35 day windows, from t_1 to t_{35} and t_{36} to t_{71} . If a pattern occurs between, for example, t_{30} and t_{40} this would not be identified. Under the rolling window approach, this would not be the case. Essentially, the rolling window starts as t_1 to t_{35} followed by t_2 to t_{36} and so on. Although this is a computationally expensive approach, it is the appropriate way to identify all patterns and avoid look-ahead bias. If this methodology were not adopted, then *future* prices would enter the smoothing process and bias the results.

Returns from the head and shoulders patterns are measured over 1, 5, 10, 20, 30 and 60 days. Note that to avoid bias, if a rolling window finishes at, for example, t_{83} we compute returns from the following day t_{84} . The return for a pattern found in a particular window, w , is

$$r_w = \ln \frac{P_{t+1+q}}{P_{t+1}} \cdot 100 \quad (2.7)$$

where q represents the holding times that are evaluated. This means that returns from successful head and shoulders patterns will be negative (as a short sale is initiated) and positive for inverse head and shoulders. The empirical work here focusses on the excess return, which is calculated simply by subtracting the 3-month gilt. This risk-free rate is continuously compounded over the equivalent holding periods.

One key point is the time taken for a head and shoulders pattern to be detected. Clearly peaks and troughs can only be identified with 'certainty' after they have formed. For example, a peak which is later determined to form on January 7 may not be identified as such until January 13. Lo et al. only trade patterns that are identified at the close of each window period, ignoring the possibility that there may be a significant lag between pattern formation and detection. Such a gap occurs as it is necessary that some significant movement in price occurs for a local maxima/minima to be recognised.

This study proposes a way to account for this by introducing the new concept of the 'trade lag'. By recording the date that a pattern was detected, as well as the date when the right shoulder was formed (with the last peak or trough in the pattern), whether patterns go 'stale' can be investigated. In evaluating a trade lag of less than and greater than five days, results can be further analysed to investigate if this is the case. This is important as, given the popularity of the head and shoulders pattern amongst traders, any gain may be quickly arbitrated away.

2.4.5 Recognising head and shoulders patterns

Section 2.3 gave an outline of the composition of head and shoulders patterns and noted their features as identified by the practitioner literature. The basic pattern was illustrated in Figure 2.1 and Figure 2.2. This section gives a more formal definition, specifying an algorithm to recognise patterns automatically in order to accurately replicate what traders see with the human eye.

In order to fully describe the process of detecting head and shoulders patterns in financial time series data it is necessary to have a more precise definition of what these patterns consist of. In particular, clarity in describing the features of a head and shoulders formation is necessary in order to produce sound pattern geometries and detection algorithms. To this end, Figure 2.12 shows an artificial example of a head and shoulders pattern. The basic formation is relatively self




Figure 2.12: Diagram of a head and shoulders patterns with extrema labelled

explanatory; it can be seen that there is a central ‘head’, denoted by the highest price point in the formation, bordered on either side by peaks representing lower prices, referred to as the left and right ‘shoulders’.

As stated previously, the head and shoulders bottom represents a mirror image of this formation. The essence of the pattern is a series of three peaks and troughs with the feature that gives the formation the name of ‘head and shoulders’ being a central peak of a greater height than those either side. For ease of reference, the peaks and troughs are labelled HS1 to HS5. As such, the left shoulder is represented by HS1, the head by HS3 and the right shoulder by HS5. These points are used as the basis for constructing pattern recognition algorithms.

Table 2.1 translates the features of a head and shoulders pattern into a more formal structure based around the points labelled HS1 to HS7. The criteria for the inverse head and shoulders (to generate a buy signal) are




Figure 2.13: Trading the head and shoulders pattern

$$\text{IHS A} \left\{ \begin{array}{l} HS1 < HS2, HS3 < HS4, HS3 < HS4, HS5 < HS4 \\ HS3 < HS1, HS3 < HS5 \\ HS1 \text{ and } HS5 \text{ within } 1.5\% \text{ of their average} \\ HS2 \text{ and } HS4 \text{ within } 1.5\% \text{ of their average} \end{array} \right.$$

Lo et al. (2000) use this general specification for detecting head and shoulders patterns. Looking at the first row, $HS1 > HS2$ states that the peak representing the left shoulder, $HS1$, should be higher than the following trough, $HS2$. $HS2 < HS3$ states that the peak representing the 'head' ($HS3$) should be higher than the previous trough ($HS2$) and so on. The following line forces the head to be higher than the shoulders either side, i.e. $HS3$ should be higher than both $HS1$ and $HS5$.

The final two lines force a degree of vertical symmetry. This is achieved by requiring the two shoulders ($HS1$ and $HS5$) and the corresponding troughs between

POINT	PRICE
Pre-HS1	Rallies. <i>HS1 forms the left shoulder.</i>
HS1-HS2	Declines.
HS2-HS3	Rallies to a level higher than HS1. <i>HS3 forms the head.</i>
HS3-HS4	Declines to a level 'near' HS2.
HS4-HS5	Rallies, but fails to reach the height of the head (HS3) <i>HS5 forms the right shoulder.</i>
Post-HS5	Declines.

Table 2.1: Key features defining head-and-shoulders patterns, applied to the idealised pattern

shoulder and head (HS2 and HS4) to be within 1.5 percent of their mean. In imposing this restriction, the patterns detected do not 'lean' with an uptrend or downtrend.

The mirror image of the inverse head and shoulders - the conventional head and shoulders pattern - is used to generate sell signals:

$$\text{HS A} \left\{ \begin{array}{l}
 HS1 > HS2, HS3 > HS2, HS3 > HS4, HS5 > HS4 \\
 HS3 > HS1, HS3 > HS5 \\
 HS1 \text{ and } HS5 \text{ within } 1.5\% \text{ of their average} \\
 HS2 \text{ and } HS4 \text{ within } 1.5\% \text{ of their average}
 \end{array} \right.$$

HS A and **IHS A** impose a set of restrictions on the local maxima and minima derived from stock price data to identify head and shoulders tops and bottoms.

2.4.6 Evaluating head and shoulders patterns

t-tests are performed to evaluate whether returns from head and shoulders induced trades are significantly different from zero. The *t*-statistic for buys (from inverse head and shoulders) is

$$\frac{\mu_b}{(\sigma_b^2/N_b)^{1/2}} \tag{2.8}$$

where μ_b and $N - b$ are the mean return and number of buy trades respectively. σ_b^2 is the variance for the buy trades. The t -statistic for the sells is

$$\frac{\mu_s}{(\sigma_s^2/N_s)^{1/2}} \quad (2.9)$$

where μ_s and $N - s$ are the mean return and number of sell trades respectively. σ_s^2 is the variance for the sell trades.

In addition, the difference in returns from head and shoulders tops versus inverse head and shoulders can be evaluated (i.e. the difference between buy and sell trades). If head and shoulders patterns contain useful information a significant difference would be expected, with the returns from tops being negative, and the returns from bottoms being positive. The t -statistic is proposed as

$$\frac{\mu_b - \mu_s}{(\sigma^2/N_b + \sigma^2/N_s)^{1/2}} \quad (2.10)$$

where μ_b and μ_s are the mean returns for buys and sells, respectively and N_b and N_s are the number of buys and sells respectively. σ^2 is the sample variance.³⁴

In addition, the percentage of successful trades is reported. This is often referred to as the 'hit rate' in the practitioner literature. It is an important statistic to present, as we would expect that if the head and shoulders trading rules do not produce useful information that the hit rate for buys and sells should be similar.

Bootstrapping

Whilst many studies have confined themselves to evaluating profits from technical trading rules on the basis of standard t -tests, this result is not entirely satisfactory. This is because for the inference drawn from these significance tests to be sound we must be convinced that returns exhibit normality, homoscedasticity and are independent. This is commonly not the case for financial time series; furthermore,

³⁴This t -statistic is similar to that employed by Brock et al. (1992).

non-stationarity is usually manifested. One way to overcome this difficulty, and be able to validate findings, is to adopt a bootstrap approach which does not make assumptions concerning the underlying distribution.

Bootstrapping is a computer intensive technique similar in nature to a Monte Carlo strategy, as it uses the original data in order to generate a large number of simulated series. These 'pseudo price-series' can then be used to assess estimates based on the original sample. Due to the possibility of constructing a confidence interval for a parameter in which either the population distribution and/or the distribution of the statistic are not known, the bootstrap has been widely applied in economics and finance. The technique has also been used to evaluate the profitability of technical trading rules (Brock et al., 1992; Mills, 1997).

The bootstrap approach adopted in this study is to generate a large number of pseudo-price series, to which the head and shoulders trading rules are then applied. To generate these series, the price series for each security are re-arranged (or 'shuffled') with replacement. This is a similar approach to that adopted by Levich and Thomas (1993), and is particularly useful as the pseudo-price series retain the same distributional properties as the original series. Therefore, we generate a series that is distributed similarly, but where the actual course of price action is randomised.

Once a set number of random series have been generated by resampling with replacement, the head and shoulders pattern detection algorithm is run over each. As with the actual price series, excess returns for holding periods of 1 to 60 days are computed. The returns from these pseudo-price series can then be compared with the actual price series. An advantage of this bootstrap approach is that the null hypothesis is relatively simple, and allows clear inference to be drawn. If the head and shoulders trading strategy does not produce useful information, we would expect that the returns generated from the random series to be indistinguishable from those of the actual series.

2.4.7 Conclusions

This section presents the methodology that is used to investigate the profitability of head and shoulders patterns. This is broadly a two step process. First, price data are smoothed to identify local peaks (maxima) and troughs (minima). Second, a geometric definition of head and shoulders patterns is applied to these peaks and troughs. Following this process, it is possible to compute mean buy and sell returns over a number of time horizons. Furthermore, the trade lag is imposed as a restriction in order to discover how quickly, after the patterns form, any profits from a head and shoulders trading rule dissipate.

2.5 Empirical Results

The purpose of this research is to examine the profitability of the head and shoulders pattern as the best example of advanced technical analysis. This section presents the empirical results of the study. Results are first shown for the sample period, from January 1 1980 to December 31 2003. Holding periods of 1, 5, 10, 20, 30 and 60 days are examined to assess the persistence of head and shoulders profits. Results are further deconstructed based on the 'trade lag', as detailed above, to account for a lag between pattern formation and detection. This addresses the question of how quickly any gains from head and shoulders patterns are absorbed by traders.

2.5.1 Summary Statistics

Table 2.2 reports the frequency count for the number of patterns detected over the entire sample, from 1980 to 2003. Panel (a) shows the results for a globally optimised bandwidth, and panel (b) for a locally optimised bandwidth. The number of patterns annually per stock is also shown, which affords a picture of how frequently a head and shoulders trading strategy triggers trades. There are roughly the same number of head and shoulders and inverse head and shoulders patterns detected over the entire sample period. Both Lo et al. (2000) and Dawson and Steeley (2003) also record an approximate 50 per cent split between head and shoulders and inverse head and shoulders patterns over their respective samples.

It is apparent that the number of patterns detected is not uniform across the years in the sample. For both globally and locally optimised bandwidths, the fewest patterns are found in 1981 and the most in 1996. There is a substantial difference between a total of 384 patterns in 1981 and 3,405 in 1996 for globally optimised bandwidth. The total patterns in these two years for locally optimised bandwidth are 481 and 4,819, respectively. This translates into a considerable variation in the

Year	(a) Globally optimised bandwidth					(b) Local optimised bandwidth				
	Patterns			Mean per stock per annum		Patterns			Mean per stock per annum	
	Total	HS	IHS	HS	IHS	Total	HS	IHS	HS	IHS
1980	404	218	186	0.6	0.5	525	294	231	0.8	0.7
1981	384	208	176	0.6	0.5	481	240	241	0.7	0.7
1982	439	236	203	0.7	0.6	695	327	368	0.9	1.1
1983	898	417	481	1.2	1.4	1268	602	666	1.7	1.9
1984	398	208	190	0.6	0.5	698	320	378	0.9	1.1
1985	769	361	408	1.0	1.2	1144	533	611	1.5	1.7
1986	1034	491	543	1.4	1.6	1627	739	888	2.1	2.5
1987	2864	1441	1423	4.1	4.1	3970	1990	1980	5.7	5.7
1988	1530	806	724	2.3	2.1	1985	1045	940	3.0	2.7
1989	2896	1439	1457	4.1	4.2	3983	2018	1965	5.8	5.6
1990	1776	931	845	2.7	2.4	2421	1238	1183	3.5	3.4
1991	2238	1153	1085	3.3	3.1	3268	1722	1546	4.9	4.4
1992	2244	1195	1049	3.4	3.0	3192	1610	1582	4.6	4.5
1993	2621	1251	1370	3.6	3.9	3748	1931	1817	5.5	5.2
1994	2413	1134	1279	3.2	3.7	3550	1489	2061	4.3	5.9
1995	3137	1465	1672	4.2	4.8	4254	2118	2136	6.1	6.1
1996	3405	1675	1730	4.8	4.9	4819	2354	2465	6.7	7.0
1997	3059	1496	1563	4.3	4.5	4277	2062	2215	5.9	6.3
1998	2196	1126	1070	3.2	3.1	2961	1536	1425	4.4	4.1
1999	2190	1121	1069	3.2	3.1	2914	1521	1393	4.3	4.0
2000	1855	967	888	2.8	2.5	2685	1256	1429	3.6	4.1
2001	2076	1182	894	3.4	2.6	3215	1799	1416	5.1	4.0
2002	1856	957	899	2.7	2.6	2924	1630	1294	4.7	3.7
2003	2181	1108	1073	3.2	3.1	3376	1676	1700	4.8	4.9

Table 2.2: Frequency counts for occurrences of head and shoulders (HS) and inverse head and shoulders (IHS) patterns in UK stock data 1980-2003. Patterns found in the largest 350 stocks by market capitalisation sorted annually. 'Mean per stock per annum' shows the number of patterns occurring per stock per annum. Panel (a) reflects a globally optimised single bandwidth and panel (b) a locally adapted bandwidth.

mean number of patterns found per annum for each security. For example, under a globally optimised bandwidth, an average of 0.5 buy signals would have been recorded for each security in 1981. This compares with 4.9 in 1996.

This is an interesting result. Murphy (1999) makes a connection between volatility and chart patterns. One possible interpretation of this finding is therefore that the number of head and shoulders patterns occurring is related to volatility. This is supported by the greatest number of patterns occurring in 1996 and 1997 as the market approached a peak, and volatility was higher than in many preceding years. It seems likely that greater volatility allows more peaks and troughs—the building blocks of all chart patterns—to form.

Table 2.2 also allows an initial comparison between smoothing using a globally optimised bandwidth compared to a locally optimised bandwidth. For global, where one bandwidth is chosen to best fit all of the prices in each window under investigation, 2,181 patterns are recorded in total. By contrast, allowing the bandwidth to be locally optimised produces 3,376 patterns. This is not a surprising result. Where one bandwidth is selected to suit all observations in a window there are more likely to be over and under-smoothed regions. When locally optimised, the bandwidth can take account of short periods of greater variability. One of the issues that this chapter addresses is the importance of the kernel smoothing methodology. It will be further investigated below whether this larger number patterns for locally optimised bandwidth leads to greater mean profitability.

2.5.2 Is the head and shoulders profitable?

Lo et al. (2000) concede that their methodology, involving comparing the unconditional empirical distribution of returns with the conditional empirical distribution (conditioned on the occurrence of technical patterns) “does not guarantee a profitable trading strategy” (p.1726). The 1-day period over which Lo et al. compare the unconditional and conditional distribution of returns is far shorter than the

Table 2.3: Head & Shoulders returns for 1980-2003 with extrema identified using a globally optimised bandwidth

Period	N		Mean π		$\pi > 0$		t -statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	22277	22586	0.0130 (2.3034)	-0.0236 (2.1438)	0.45	0.56	0.8393 (0.4013)	-1.6515 (0.0987)	1.7381 (0.0822)
5	22277	22586	-0.0078 (4.4713)	-0.0524 (4.7032)	0.50	0.51	-0.2604 (0.7946)	-1.6726 (0.0944)	1.0285 (0.3037)
10	22277	22586	-0.0535 (6.3217)	-0.0509 (6.6606)	0.51	0.50	-1.2607 (0.2074)	-1.1465 (0.2516)	-0.0424 (0.9662)
20	22277	22586	-0.2844 (9.1719)	-0.1025 (9.2007)	0.51	0.49	-4.6154 (0.0000)	-1.6699 (0.0950)	-2.0975 (0.0360)
30	22277	22586	-0.4716 (11.9001)	-0.1826 (11.1796)	0.52	0.47	-5.8934 (0.0000)	-2.4463 (0.0144)	-2.6507 (0.0080)
60	22277	22586	-0.4385 (17.0075)	-0.4681 (16.4762)	0.53	0.46	-3.8147 (0.0001)	-4.2407 (0.0000)	0.1875 (0.8513)

The sample period is January 1 1980 to December 31 2003 comprising the 350 largest stocks (resampled annually) by market capitalisation. Dead stocks are included. "Period" is the holding period, i.e. 30 would represent the return from t_1 to t_{30} where t_1 is the buy date. "N Buy(Sell)" represents the number of buys(sells). "Mean π Buy (Sell)" is the mean return for buys(sells), with standard deviation shown below in parentheses. " $\pi > 0$ " shows the percentage of profitable trades for buys and sells (the 'hit rate'). t -statistics are shown with p -values below in parentheses. Buy-Sell reflects a standard t -ratio for the difference between mean buy and sell trade returns. For ease of reference, these results are reported as percentages (i.e. -0.4681 for the 60 day sell holding period is simply -0.4681%).

holding periods employed by traders. Indeed, this is a key shortcoming of their “natural first step in a quantitative assessment of technical analysis”. Establishing the profitability of a technical trading strategy based on the head and shoulders is a central research question in this chapter. To address this, the actual profitability of a trading strategy based upon visual price patterns—represented by the head and shoulders—is investigated. The results here look at 1-day, 5-day, 10-day, 20-day, 30-day and 60-day returns from inverted head and shoulders buys and head and shoulders sells.

Table 2.3 reports the mean profitability of head and shoulders trades using a globally optimised bandwidth. Results are separated into holding periods of 1, 5, 10, 20, 30 and 60 trading days as shown in column 1. Columns 2 and 3 show the number of head and shoulders and inverse head and shoulders patterns detected during the period 1980 to 2003. For ease of reading, head and shoulders patterns are labelled ‘sell’ and inverse head and shoulders patterns are labelled ‘buy’. In this instance, there are 22,277 buys and 22,586 sells over the entire sample period. ‘Mean π ’ shows the mean profitability of buy and sell trades derived from inverse and non-inverse head and shoulders patterns, respectively. The standard deviation is given in parentheses beneath the mean return.

On first inspecting the mean buy and sell returns in Table 2.3, the most striking result is that, with the exception of 1-day, all the other holding periods exhibit negative excess buy returns. All of the mean sell returns are negative. However, for holding periods of 10, 20 and 30 days, the negative mean buy return is actually greater in magnitude than the mean sell return. Clearly, this does not appear to be the basis of a profitable trading strategy. The mean excess returns increase in magnitude as the holding period increases from 1 to 60-days. One possible reason for this is that if the inverse head and shoulders is not informative then, in a generally rising market, a longer holding period would in and of itself show higher returns. Therefore, a longer holding period would demonstrate higher returns

regardless of whether there is useful information conveyed by head and shoulders patterns or not.

To better explain these returns it is useful to look at the percentage of trades which generated returns in the expected direction (positive for buys and negative for sells). The column labelled ' $\pi > 0$ ' in Table 2.3 shows this for buy and sell trades. This statistic is sometimes referred to as the 'hit rate' or 'success rate' in the academic as well as practitioner. If a trading strategy based on the head and shoulders pattern does lead to profitable signals, the hit rate should be the same. This is generally the case, with a hit rate ranging between 45% and 56%.

The final three columns in the table present t -statistics. These are first presented for buy trades against a null of a mean return of zero. The p-value is computed and is shown in parentheses beneath each t -statistic. With reference to the p-values, it can be seen that the null can be rejected for the longer holding periods of 20, 30 and 60 days. At these horizons the buy returns are significantly different from zero. This is not the case for the 1, 5 and 10 day holding periods. For the sell returns, the null of returns being equal to zero can be rejected for only 30 and 60 days.

It is slightly more interesting to look at the t -statistic computed to test the difference between buy and sell returns 'Buy-Sell' (see the previous section for details on its construction). Looking to the p-values it, can be seen that for holding periods of 20 and 30 days there is a significant difference between buy and sell returns.

If the head and shoulders pattern provided useful trading signals, we would expect that the mean buy return to be positive, the mean sell return to be negative, and these returns to be significant. This is not the case. Although the sell returns produce a mean return in the expected direction (in contrast to the mean buy returns) these prove to be insignificant for all but two cases. In one case, for a holding period of 30 days, the negative mean buy return is (significantly) greater in magnitude than the mean sell return. Whilst the mean sell return for 60 days is

significant, it is only slightly greater in magnitude than the negative buy return, and there is no significant difference between the two. Overall, this table appears to confirm that head and shoulders patterns do not provide economically useful information. This is an important result given the emphasis that traders place on chart patterns, and the head and shoulders in particular.

Given the central role that detecting peaks and troughs in noisy price data has in chart pattern recognition, one of the aims of this chapter is to evaluate how useful kernel regression is for this purpose. It was explained above that the Gasser-Muller kernel estimator with a globally optimised bandwidth is a more attractive methodology than the Nadaraya-Watson kernel estimator with cross-validation. The latter is adopted by Lo et al. (2000), and it has also been noted that their approach of taking 30% of the cross-validated bandwidth is arbitrary. In order to verify this, we first calculate the globally optimised bandwidth and also take 30% of this, before smoothing the price data.

Table 2.4: Head & Shoulders returns for 1980-2003 with extrema detected using a 30% of the globally optimised bandwidth ($h \times 0.3$).

Period	N		Mean π		$\pi > 0$		t -statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	55497	55935	-0.0100 (2.6191)	-0.0135 (2.3669)	0.45	0.55	-0.8972 (0.3696)	-1.3522 (0.1763)	0.2379 (0.8120)
5	55497	55935	-0.0563 (4.9169)	-0.0265 (4.9001)	0.50	0.50	-2.6939 (0.0071)	-1.2766 (0.2017)	-1.0136 (0.3108)
10	55497	55935	-0.1364 (6.8365)	-0.0648 (6.9225)	0.50	0.50	-4.6907 (0.0000)	-2.2115 (0.0270)	-1.7361 (0.0825)
20	55497	55935	-0.2663 (10.0877)	-0.1004 (9.6237)	0.50	0.48	-6.1907 (0.0000)	-2.4578 (0.0140)	-2.8090 (0.0050)
30	55497	55935	-0.3100 (12.5839)	-0.1112 (11.6412)	0.51	0.47	-5.7656 (0.0000)	-2.2480 (0.0246)	-2.7369 (0.0062)
60	55497	55935	-0.5334 (17.9664)	-0.3556 (16.6312)	0.52	0.46	-6.9070 (0.0000)	-5.0131 (0.0000)	-1.7144 (0.0865)

Table 2.4 reports the results from modifying the optimal bandwidth in this manner. Reducing the bandwidth in kernel smoothing produces a more ‘wiggly’

line, tracking the original series more closely. It is therefore expected that more peaks and troughs will be detected as the derivative of this series possesses more turning points. The table shows that this is the case: approximately 55,000 buys and sells are generated with the smaller bandwidth, compared to approximately 22,000 using the globally optimised bandwidth.

Whilst reducing the bandwidth in a similar manner to Lo et al. (2000) allows more patterns to be detected, and hence more buy and sell trades, the table shows little evidence that this results in a more successful trading strategy. In fact, it appears to be less profitable than previously. In particular, all of the mean buy and sell returns are now negative. The fraction of successful trades is broadly similar, and there is only a significant difference between mean buy and sell returns in two cases. Furthermore, because of the much greater number of trades in this case, the transactions costs of this strategy would be greater.

Table 2.5: Head & Shoulders returns for 1980-2003 with extrema detected under global bandwidth optimisation with $h \times 2$.

Period	N		Mean π		$\pi > 0$		<i>t</i> -statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	2275	2530	0.0687 (3.6393)	-0.0486 (2.4750)	0.45	0.61	0.9009 (0.3677)	-0.9882 (0.3231)	1.3181 (0.1875)
5	2275	2530	0.1883 (5.1564)	-0.1255 (4.5504)	0.51	0.56	1.7413 (0.0818)	-1.3874 (0.1654)	2.2404 (0.0251)
10	2275	2530	0.3080 (6.8661)	-0.1084 (6.2791)	0.52	0.56	2.1383 (0.0326)	-0.8675 (0.3858)	2.1947 (0.0282)
20	2275	2530	-0.0079 (9.3915)	-0.3759 (9.2921)	0.52	0.54	-0.0401 (0.9680)	-2.0334 (0.0421)	1.3635 (0.1728)
30	2275	2530	-1.0125 (14.6138)	-0.9265 (11.4180)	0.50	0.55	-3.2981 (0.0010)	-4.0789 (0.0000)	-0.2286 (0.8192)
60	2275	2530	-1.7003 (19.1244)	-1.2406 (17.6607)	0.52	0.54	-4.2089 (0.0000)	-3.5237 (0.0004)	-0.8664 (0.3863)

This result is important in showing that making an arbitrary adjustment to the bandwidth does not translate into a more profitable trading strategy. To further investigate this issue, Table 2.5 presents results from multiplying the globally

optimised bandwidth by two. With this greater smoothing, there are far fewer extrema identified, with only 2,275 and 2,530 buy and sell trades, respectively.

The results from doubling the globally optimised bandwidth are interesting. In contrast to previous findings, the mean buy returns for holding periods of 1, 5 and 10 days are now positive. The size of these returns generated by the smaller number of patterns is also greater. Together, this suggests that increasing the optimised bandwidth is more appropriate than reducing it, counter to the conclusion of Lo et al. (2000).

However, the results presented in Table 2.5 are still unsupportive of the head and shoulders. In particular, because only two of the buy-sell returns are significant. Further, the 1, 5 and 10 day buy returns with a positive sign (unlike previous results) are insignificant.

Given that altering bandwidth to change the nature of the extrema detected does not alter the conclusions, it seems prudent to adopt the optimised bandwidth, which allows a consistent approach. In any case, the question of how to adjust the optimised bandwidth, even as to whether to increase or decrease it, is arbitrary and presents a clear risk of data mining. Up to now, the global optimum bandwidth has been used. The previous section noted that we can also use a locally optimised bandwidth. This has several attractions, most notably that it is more likely to be able to take account of time varying volatility, which is present in almost all financial time series.

Table 2.6 presents results from detecting extrema using the locally optimised bandwidth methodology proposed by Herrmann (1997). It is most appropriate to compare these findings with those from the globally optimised bandwidth shown in Table 2.3. First, the number of patterns found is larger with a locally optimised bandwidth; there are around 32,000 buys and sells compared to around 22,000 previously. It is expected that local optimisation is advantageous as bandwidth is data driven, and can increase or reduce as needed to obtain the best fit. Note that

Table 2.6: Head & Shoulders returns for 1980-2003 with extrema detected using a locally optimised bandwidth.

Period	N		Mean π		$\pi > 0$		<i>t</i> -statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	31930	32050	0.0096 (2.3416)	-0.0244 (2.1054)	0.45	0.55	0.7359 (0.4618)	-2.0784 (0.0377)	1.9362 (0.0528)
5	31930	32050	0.0017 (4.4968)	-0.0495 (4.6206)	0.50	0.50	0.0689 (0.9451)	-1.9168 (0.0553)	1.4212 (0.1553)
10	31930	32050	-0.0737 (6.3416)	-0.0332 (6.4731)	0.50	0.49	-2.0726 (0.0382)	-0.9173 (0.3590)	-0.7984 (0.4246)
20	31930	32050	-0.3320 (9.2994)	-0.0097 (8.9666)	0.50	0.48	-6.3601 (0.0000)	-0.1936 (0.8465)	-4.4615 (0.0000)
30	31930	32050	-0.5824 (11.7193)	-0.0124 (10.9460)	0.51	0.46	-8.8455 (0.0000)	-0.2027 (0.8394)	-6.3552 (0.0000)
60	31930	32050	-0.4399 (16.6902)	-0.4937 (16.3076)	0.53	0.46	-4.6661 (0.0000)	-5.3763 (0.0000)	0.4120 (0.6803)

this fitting approach does not mean look-ahead bias is introduced. This is because of the rolling window method employed.

The results show that this approach does not produce a much improved head and shoulders trading strategy. In general, the results from using a locally optimised bandwidth are quite similar to the global optimum. There is a slight improvement in the gap between mean buy and sell returns in some cases, although buy-sell remains insignificant apart from holding periods of 20 and 30 days (and little importance can be attached to this given that the mean buy returns are negative). Although there is not a marked difference between global or local optimisation, the latter technique can still be considered preferable. In particular, this is appealing given that bandwidth can increase or reduce in line with the volatility in the data.

The above results have shown that a trading strategy based around the head and shoulders chart pattern does not appear to be profitable. Whilst a small number of mean buy and sell returns are significant and in the direction predicted by the pattern, the findings are inconsistent. Such a strategy could therefore not

Table 2.7: Head & Shoulders returns for 1980-2003 with a trade lag of ≤ 5 and extrema detected using a locally optimised bandwidth.

Period	N		Mean π		$\pi > 0$		t-statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	9163	9467	0.0579 (2.9396)	-0.0930 (1.9869)	0.46	0.59	1.8849 (0.0595)	-4.5548 (0.0000)	4.1149 (0.0000)
5	9163	9467	0.2039 (4.6937)	-0.3498 (4.8315)	0.51	0.54	4.1547 (0.0000)	-7.0391 (0.0000)	7.9180 (0.0000)
10	9163	9467	0.0974 (6.5855)	-0.5319 (6.9614)	0.51	0.53	1.4120 (0.1580)	-7.4204 (0.0000)	6.3277 (0.0000)
20	9163	9467	-0.0218 (9.0153)	-0.4254 (9.2704)	0.51	0.52	-0.2302 (0.8180)	-4.4514 (0.0000)	3.0107 (0.0026)
30	9163	9467	-0.2446 (11.6354)	-0.4230 (10.8723)	0.52	0.49	-2.0016 (0.0454)	-3.7665 (0.0002)	1.0819 (0.2793)
60	9163	9467	-0.2393 (17.2894)	-0.7792 (16.1643)	0.53	0.47	-1.3097 (0.1903)	-4.6502 (0.0000)	2.2024 (0.0276)

be profitably employed by traders. However, one of the questions posed in this chapter is to investigate whether the time between the formation of a particular pattern and when it is detected by kernel smoothing is important. To address this, Table 2.7 shows results from again using a locally optimised bandwidth, but imposes the restriction that there must be a gap of less than five trading days between the last peak or trough completing the pattern, and the time at which it is detected. With a five day trade lag, buy and sell signals are only traded on if the pattern is found within five days of its completion - these are therefore ‘fresher’ patterns. This leads to a research question of whether such patterns produce better buy and sell signals. This approach is new; previous research has ignored this aspect making it a valuable addition to this study.

In terms of the frequency of head and shoulders patterns, Table 2.7 shows that imposing the restriction that a pattern should have completed within five days prior to its recognition considerably reduces the number of trades. There are 9,163 buy trades and 9,467 sell trades compared with 31,930 and 32,050 buy and sell trades with no restriction, as shown in Table 2.6. In this unrestricted case, the mean

buy returns over holding periods of 10, 20, 30 and 60 days were negative. Now, with the restriction in place, the signs of the 1, 5 and 10 day mean buy returns are positive, as predicted by the inverse head and shoulders pattern. The magnitude of these returns is also far greater. For instance, the significant excess buy return for a 5 day holding period is approximately 10.2% per cent annually, compared with just 0.085% where no trade lag restriction is imposed.

The size of all of the mean sell returns is greater too and, compared to when there is no trade lag restriction are significant. The difference between buy and sell returns is significant in all cases with the exception of a holding period of 30 days. Taking these results overall, the length of delay between identifying a pattern, and trading on it, appears crucial in terms of capturing profits from buy trades. Furthermore, negative mean buy returns with the five day trade lag restriction for 20, 30 and 60 days suggest that most of the profits occur soon after formation. In the case of the buy trades resulting from inverse head and shoulders patterns, the greatest mean excess return is when holding for five days. This is not the case for sell trades from head and shoulders tops, where the highest excess return is seen at 60 days.

To clarify this important result, and further investigate how important the time between detection and trading is, Table 2.8 displays mean excess returns for a trade lag of between 5 and 10 days. We are therefore comparing results for patterns that completed slightly longer in the past. The results show that a similar number of patterns are found. However, evidence of profitability is not compelling compared to the results for a trade lag of ≤ 5 days. Therefore, the time between pattern occurrence and the point at which it is detected and able to be traded is critical.

Previous results had a broadly similar number of head and shoulders and inverse head and shoulders patterns and thus similar numbers of buy and sell trades. Table 2.7 shows that there are about three times more buy trades than sell trades. Imposing a trade lag restriction of five days means that many of the sell

trades seen previously are filtered out.

Table 2.8: Head & Shoulders returns for 1980-2003 with a trade lag of > 5 and ≤ 10 and extrema detected using a locally optimised bandwidth.

Period	N		Mean π		$\pi > 0$		<i>t</i> -statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	8491	8709	-0.0146 (1.9467)	-0.0119 (2.4827)	0.45	0.57	-0.6914 (0.4893)	-0.4460 (0.6556)	-0.0805 (0.9358)
5	8491	8709	-0.1009 (4.1640)	-0.0720 (4.8551)	0.49	0.51	-2.2302 (0.0258)	-1.3824 (0.1669)	-0.4186 (0.6755)
10	8491	8709	-0.1773 (5.7921)	0.0745 (6.4169)	0.49	0.50	-2.8163 (0.0049)	1.0822 (0.2792)	-2.6991 (0.0070)
20	8491	8709	-0.3625 (8.6753)	0.1531 (8.7059)	0.50	0.48	-3.8362 (0.0001)	1.6366 (0.1018)	-3.8887 (0.0001)
30	8491	8709	-0.8062 (11.5902)	0.0761 (11.0223)	0.50	0.46	-6.3831 (0.0000)	0.6413 (0.5213)	-5.1126 (0.0000)
60	8491	8709	-0.4415 (16.2872)	-0.4810 (16.3106)	0.53	0.46	-2.4747 (0.0134)	-2.7290 (0.0064)	0.1590 (0.8737)

2.5.3 Transaction and short selling costs

As noted in the review of the literature, previous studies do not adopt a consistent approach to adjustment for transaction costs. For instance, as Lo et al. (2000) specifically do not set out to examine profitability but rather compare the unconditional and conditional one-day returns, they do not consider trading costs. Savin et al. (2007) do, however, consider one-way break-even costs in relation to raw excess returns, noting figures of 0.18% (Jones, 2002) and 0.23% (Berkowitz et al., 1988) for an institutional trader. A similar approach can be taken here in order to assess the impact of transactions costs on the returns to head and shoulders patterns.³⁵

Table 2.7, discussed above, displays the excess returns from head and shoulders patterns detected using locally optimised bandwidth with a trade lag of five days. The one-way break-even transaction cost is half of the excess return. For instance,

³⁵This methodology gives an approximation of the effect of transaction costs. However, further research could consider computing the return for each trade, less transaction costs, to arrive at a figure for mean excess returns in the presence of costs.

Table 2.9: Bootstrap results from 500 simulated series compared to the actual price series.

Holding Period	Fraction of simulations greater than actual series			
	Buy	Sell	σ_b	σ_s
1	0.0	0.0	99.8	100.0
5	0.0	0.0	0.0	0.0
10	100.0	0.2	0.0	0.0
20	100.0	0.0	0.0	0.0
30	99.8	1.0	100.0	100.0
60	100.0	46.4	100.0	100.0

As detailed in the methodology section, the original price series is ‘shuffled’ (resampled with replacement) 500 times. The algorithm for identifying head and shoulders patterns is run on these pseudo price-series. For the mean, buy, sell and standard deviations of buy and sell returns (σ_b and σ_s , respectively), the columns report the fraction of simulations greater than the original series. Results are presented for 1 to 60 day holding periods.

in the case of the 60-day sell return the break-even cost is 0.3896%. The excess return is therefore greater than the estimated institutional costs noted above of between 0.18% and 0.23%. With the aim of being conservative and therefore taking the higher figure of 0.23%, excess returns for sell trades remain negative (the expected direction as these are short sales) in the case of 10-days and 60-days. Transaction costs for buy trades of 20, 30 and 60 days are irrelevant as these strategies do not show profits. However, for 1, 5 and 10 days there is a one-way break-even cost of 0.0290%, 0.1020% and 0.0487%, respectively. This implies that the some of the profits from this strategy may be subsumed by transactions costs. However, liquidity traders may be able to increase returns by utilising the information contained in head and shoulders patterns.

It is also necessary to note that short selling costs and constraints may also be relevant here. In many markets, there are legal constraints on short selling such as the uptick rule imposed by the NYSE and AMEX in the US and, in the UK, unit

trusts are prohibited from short selling activities. The costs of short selling include loan fees as well as some degree of risk that the borrowed stock will be recalled before the borrower wishes to close their trade.³⁶ For the UK, Mackinson Cowell (2005) estimate stock lending fees to be, on average, around 0.14% of the total loan value, per annum. However, these are said to range between 0.05% and 4% or above. In the case of this study, which uses data for large UK stocks, it is expected that lending fees will be at the lower end of this range.

It is important to note that traders employing short sales as part of a head and shoulders trading strategy could use methods other than borrowing stock to sell short in order to profit from downward price moves. In particular, it would be possible to purchase put options and utilise single stock futures.³⁷ However, the costs may still exceed long trades. Therefore, it would be desirable for future research in this area to directly investigate short sales costs as part of such a technical trading strategy.

2.5.4 Bootstrap tests

The results above allow interesting conclusions to be drawn concerning the profitability of advanced technical analysis. The methodology for identifying head and shoulders proposed by Lo et al. (2000) is significantly developed into a clear trading strategy to establish whether profits are present for UK securities over a range of time horizons. The finding that useful information is provided by the 'freshest' patterns, with the opposite being the case for 'staler' patterns, were evaluated with *t*-statistics. This is potentially problematic given that most financial time series are non-stationary and that we cannot assume a normal distribution. One way to resolve this issue is to use bootstrapping.

Table 2.9 presents results from a bootstrap analysis of (I)HSA, where local maxima and minima are found by a locally optimised bandwidth. The results from

³⁶Thomas (2006) provides useful detail and references on the nature of short selling costs.

³⁷Individual investors may also use contracts for difference or spread bets.




Figure 2.14: Cumulative mean standard deviations over 500 simulations.

the original series were presented in Table 2.6. The results of the bootstrapping first show the percentage of the simulated series that have mean buy return, buy standard deviation, mean sell return and sell standard deviation that are greater than the actual series. 500 replications were performed. To give this number perspective, Figure 2.14 shows how the standard deviations converge as the number of pseudo-series increases. Similar to Brock et al. (1992), the estimates settle relatively quickly.

Looking at the results for buy trades, for holding periods of 1 and 5 days, none of the simulated series produced larger mean returns than the original price series. These results can be thought of as simulated p -values.³⁸ This allows significance to be attached to the mean excess buy returns of 0.0096% and 0.0017% for holding periods of 1 and 5 days, respectively, shown in Table 2.6. On the other hand, all of the simulated series returns for 20 and 60 days, as well as all but one at 30 days, are larger than the 'real' series. This is unsurprising because, as discussed previously, all of these returns were negative.

A slightly different picture is presented for the sell trades. In this case, almost all of the simulated series fail to present returns as low as the original price series, with the exception of the 60 day sell return. For the 20 and 30 day holding periods, these results agree with the standard t -statistics presented earlier. However, in the results shown in Table 2.6, the 1 and 5 day sell returns were insignificant (albeit the 1-day return p -value is 5.28%). Brock et al. (1992) observed similar results in their bootstrap analysis, and determined that this "suggests that the distributional assumptions of the standard tests may have an impact on statistical inferences" (p.1749). This appears to be the case here. Given this, the value of performing the bootstrap simulation is clearly demonstrated.

Standard deviations of the buy and sell returns from the simulated series compared to the original are also shown in the table. The σ_b shows that all of the

³⁸This is a similar approach to Brock et al. (1992) and Mills (1997).

standard deviations from the pseudo price series were greater than the original series for returns over 30 and 60 days, and nearly all (99.8%) for 1 day. By contrast, none of the simulated series recorded a standard deviation greater than the actual series for returns over 5, 10 and 20 days. Therefore, at these time horizons, the inverse head and shoulders patterns seem to identify trades with lower volatility than would be expected by chance. Sell standard deviations (σ_s) show the same pattern.

Taken together, these results are very interesting. We are now confident that the head and shoulders pattern does appear to provide useful information. Specifically, for buy trades over a period of up to 5 days, and for sell trades up to 30 days. In the latter case, this persistence of excess returns is definitely not consistent with an efficient market. However, this would not be the case if traders were simply being rewarded for bearing increased risk. The standard deviation bootstrap results are important in addressing this question. The finding is that at 5, 10 and 20 days, head and shoulders and inverse head and shoulders appear to have *lower* return volatility than would be expected if head and shoulders patterns could produce no useful information. Accordingly, we can infer that the significant excess returns for buy trades at 5 days and sell trades at 5, 10 and 20 days are not because these trades are riskier.

2.6 Conclusions

The empirical work in this chapter has extended and moved a significant distance beyond the “first step” taken by Lo et al. (2000). Using a large dataset of UK stocks from 1980 to 2003, kernel smoothing was employed to detect peaks and troughs in stock price series. Using these localised maxima and minima, head and shoulders patterns were identified after definition of a suitable pattern geometry. Mean excess buy and sell returns were computed and evaluated in the context of a trading strategy. Unlike previous work, buy and sell returns were shown for a variety of holding periods up to 60 days. Furthermore, a new methodology is developed to investigate whether the information contained in head and shoulders patterns is quickly impounded into prices, using the new concept of a trade lag.

Several important methodological advances have been made. First, given doubt about subjective alterations to the bandwidth used in kernel regression to smooth price data, alternative approaches were investigated. In particular, more recent approaches using direct plug-in bandwidths were adopted. Second, given the apparent theoretical and practical advantages of locally optimising the bandwidth, this was compared to global optimisation. Third, the trade lag was developed and implemented, allowing the nature of returns from head and shoulders patterns to be investigated.

Results presented in the previous section proved to be interesting. In broad terms, the head and shoulders pattern appears to provide some useful information. Analysis and comparison of modifications to bandwidth, in the manner chosen by previous work, demonstrated that this appeared to be an unsound approach. To address this, results from local optimisation of bandwidth to detect peaks and troughs were evaluated. The results showed that the different approach to smoothing resulted in identification of a greater number of head and shoulders patterns. Mean excess sell returns exhibited significant profitability over 20-60 days, with equivalent annual excess returns in the order of 2% for the longest

holding period. Although considerably less at shorter holding periods, these returns suggest that traders holding positions for between 30-60 days (similar to the period over which the patterns formed) could profitably employ this strategy in the presence of moderate transactions costs. However, one important finding was that buy trades, from head and shoulders bottoms, performed less well. In fact, negative excess returns were seen from 10-60 days after identification of the chart pattern.

Given that previous work has not established whether the time between completion of the chart patterns and their identification is important, the trade lag was introduced. When set to ≤ 5 the number of patterns identified diminished considerably. A clearly identifiable change in the magnitude of excess returns was seen as a result. For instance, with no trade lag the mean excess return from selling short and covering the trade 30 days later was just 0.10% on an annual basis. With the restriction that patterns must have formed less than five trading days ago, the mean excess return becomes 3.5%. These results are highly significant: the economic use of the head and shoulders is clearly linked to how quickly patterns are identified and traded upon.

Given that these results present a new and different picture than previous work, bootstrapping was carried out to verify significance. The results supported the earlier analysis and, in addition, gave an insight into the riskiness of returns from head and shoulders patterns. Based on these results, the head and shoulders actually appears to identify trades with low volatility of returns. However, although this result suggests that returns are not a reward for bearing additional risk, it would be desirable for future research to look at risk-adjusted excess returns using the standard three-factor model. In addition, given that short-term trends play a role in the formation of head and shoulders patterns, it would also be desirable to augment the three-factor model with a momentum factor when looking at risk-adjusted returns.

Chapter 3

Removing the Straight-Jacket on Technical Analysis

3.1 Introduction

The first empirical chapter in this thesis made a major contribution to the study of technical analysis on a number of fronts. First, by looking at the profitability of trading rules based upon the head and shoulders pattern, our knowledge of advanced technical analysis was enhanced. Profitability was scrutinised over a number of time horizons from 1 to 60 trading days. Second, the new concept of the 'trade lag' was introduced to investigate if returns to head and shoulders patterns were affected by the time elapsed between the completion of a head and shoulders chart pattern and the time at which it could be positively identified, and therefore traded upon. This is of particular importance to traders, especially if they use an automated (or systematic) trading system to identify chart patterns. The chapter was supported by the use of a large dataset of UK stocks for the period January 1 1980 to December 31 2003. This chapter builds upon and extends this work in a number of important ways.

The previous results were valuable for developing and evaluating a trading strategy based upon the proposal for further research put forward by Lo et al. (2000), and as such constitute a distinct and valuable 'second step' in a rigorous empirical study of technical analysis. In this chapter, it is argued that the restrictions imposed by the limited previous studies into visual chart patterns are potentially problematic. With the knowledge that technical analysis is very heavily applied in financial markets, it is desirable that the actions of traders should be replicated as closely as possible in evaluating technical trading strategies. This chapter demonstrates that patterns detected using the criteria developed by Lo et al. and related studies would, in reality, be unrecognisable by professional technical analysts looking for head and shoulders formations. Conversely, further patterns that would provide trading signals to professionals may not be detected at all. The key aim of this chapter is to address this important issue.

The first additional contribution of this chapter relates to the importance of

closely replicating professional traders in gaining an understanding of the profitability of the head and shoulders pattern as emphasised above. This chapter presents an important advancement in developing and evaluating pattern geometries that are different from those investigated previously, and identifiable with those employed by traders. We separately analyse the introduction of the neckline and separate continuation from reversal patterns by looking at the prevailing trend.

Given that one of the goals of this chapter is to more closely align research on technical analysis with the activities of professional technical analysts in financial markets, a detailed study of the practitioner literature is undertaken. This is composed of numerous monographs and other texts, which have been afforded a wide audience among technical analysts. This is an important element, as the practitioner literature is largely ignored in the limited amount of previous work. In doing so, the patterns identified are far more closely associated with those recognised by professional technical analysts.

One issue that arises from looking at how technical analysts operate is that the head and shoulders pattern is often perceived to form over a longer period than that previously evaluated. To address this, holding periods of 35 and 65 days are examined. These are further evaluated using the trade lag concept developed in Chapter 2.

The robustness of results is evaluated using bootstrapping. The approach of simulating pseudo-price series as a random walk is used to benchmark the results from the actual prices series. This important aspect is ignored in previous work, notably Savin et al. (2007).

These contributions will provide a significant improvement in our understanding of advanced technical analysis. Crucially, in contrast to previous work, this study closely aligns empirical work with the actual activities of professional traders. In doing so, the results are not only valuable to academics, but to traders as well.

3.2 Organisation

This chapter is divided into three sections. First, Section 3.3 provides a review of the practitioner literature. This provides the central framework for establishing a geometric specification for head and shoulders patterns that practitioners use. Second, the data and methodology section builds upon that presented in the previous chapter; given the practitioner literature, new trading rules for the head and shoulders are developed.

Finally, empirical results are presented and discussed in Section 3.5, evaluating the profitability of the head and shoulders pattern, following the analytical and methodological developments made in this chapter. Conclusions to the study are presented in Section 3.6.

3.3 Practitioner literature

The 'practitioner literature' encompasses a wide range of publications, connected by their focus for consumption by traders. Writing on technical analysis can be traced back at least until the start of the twentieth century. Given the assertion that previous academic inquiry has failed to capture patterns that traders would recognise, it is vital to address this. Close study of the large amount of literature makes this possible. This section presents a review of the most important elements of the practitioner literature pertaining to head and shoulders patterns. In particular, pursuant to the aims of this chapter, contributory evidence supporting a more realistic set of definitions of the head and shoulders pattern is presented.

3.3.1 The history of technical analysis

Largely perceived as the 'father' of technical analysis, Charles Dow wrote prolifically on the subject in a series of articles published in the Wall Street Journal around 1900. While Dow's prominence is largely due to his creation of market indices, his ideas on how information was compounded into prices and trading strategies have a clearly identifiable impact today. These articles constitute what we now refer to as "Dow's Theory", a term coined by Nelson (1903) who collected Dow's idea in a book published the year after his death.

As well as an early proposal that prices impound all available information in the manner of an efficient market, Dow developed ideas that still underpin technical analysis today. First, he determined that prices trended, with three main categories of trend: primary, secondary and minor. He observed that while prices often moved against the direction of the primary or secondary trend, more often than not they would revert to them in due course.¹

¹Dow identified three main phases in trends based around accumulation by investors, buying by the public who are following trends, the 'public participation' phase and finally the 'distribution phase' where the greatest changes in price occur and those investors who initially bought sell and distribute their stock to others. A comprehensive discussion can be found in Murphy (1999, p.26)

Dow identified that the multiple indices he created could be used together to give a more informed forecast of market direction. He further established that volume plays a role in confirming trends. A trend formed on low volume was not given as much significance as one formed with high volume. The ideas that Dow proposed about trend formation were accompanied with detailed discussion of how trends reverse. Particularly relevant for this study, Dow developed patterns in price movements that could be used to identify trend reversals. These patterns bear a remarkable resemblance to head and shoulders patterns, which also seek to identify reversals in price movements.²

Schabacker (c1932) developed Dow's ideas and formulated a system for their use in trading individual stocks.³ With reference to reversal formations, continuation patterns, trend lines and details of support and resistance, this text lays the foundations of modern technical analysis.

While Dow is often considered the 'father' of technical analysis, there is an earlier example of its use. Rice traders in the 17th century have been shown to have used so-called 'candlestick' charts (Nison, 2001). The candlestick approach involves plotting the open, close, high and low price on a chart with the view that observing this would lead to more informed trading decisions. However, the first 'conventional' writings on technical analysis can be attributed to Dow.

King (1934) provides a summary of an *American Statistical Association* conference on "Technical Methods of Forecasting Stock Prices". This shows how quickly technical theories developed. Head and shoulders patterns are discussed in their role of signalling a trend reversal. This is important, as a long history of use has led to a clear idea of the head and shoulders pattern, reducing scope for data mining.

William Hamilton—like Dow, an editor of *Wall Street Journal*—developed Dow's ideas further. In a series of editorials, he extended his mentor's thoughts on market

and Edwards and Magee (2001).

²What Dow refers to as 'lines' constitute sideways patterns in prices akin to what we would now view as rectangle formations.

³Harriman House re-published the work recently (Schabacker, 2005).

averages and arrived at the notion that the movement of price averages could 'confirm' each other. Hamilton's ideas were subsequently published in a short book (Hamilton, 1922). Following this, Dow Theory was extended and refined by several authors (Rhea, 1932; Schaefer, 1956). However, the essence remained and carries through to technical analysis today.

Reviewing the impact of the 1929 depression and subsequent bull market on American stocks, Gann (1936) published a work entitled "New Stock Trend Detector". This was very much a practitioners' text, with rules to detect the trends and turning points in prices. Gann supports this with the results of trades in Chrysler stock made in the preceding ten years according to his rules.

These early proponents laid the groundwork for what we still understand as technical analysis today. The concepts of support and resistance, averages and deriving indicators from price movements are all actively used by traders. However, whilst relevant to the investigation of the subject of technical analysis as a whole, these concepts can largely be understood to be 'basic' technical analysis, and as such have been widely examined in the academic literature discussed in Chapter 2. Head and shoulders patterns are the main subject of investigation in this work, and as such it is valuable to ascertain when they were first identified in the practitioner literature.

Head and shoulders patterns depend upon the concepts of support and resistance and reversals from a prevailing trend, as discovered in the early practitioner literature. However, it was not until 1948 that what we would understand as a head and shoulders pattern today was presented in *Technical Analysis of Stock Trends*. Now in its 8th edition after selling some 850,000 copies, this work can be viewed as one of the definitive references on technical analysis. Given the publication date of the first edition, there is some justification in asking why—if we assume an efficient market—profits from these patterns have not disappeared, and, if this is not the case, why there has been a lack of academic investigation into the subject.

Feature	Characteristic
Left shoulder	Strong rally providing the climax to an extensive advance
The 'head'	Advance reaching a peak higher than that seen previously in the left shoulder, followed by a retracement to a level close to that of the previous retracement
Right shoulder	Another rally, failing to reach the height of the head
Confirmation	Decline in prices below the 'neckline', which is drawn from the points of the troughs on either side of the head.

Table 3.1: Key features defining head-and-shoulders patterns, summarised from Edwards and Magee (2001, p.57). Whilst these represent the head and shoulders top pattern (for initiating a short sale), the characteristics for a head and shoulders bottom to initiate a long trade are analogous.

3.3.2 The head and shoulders pattern

The practitioner texts show the defining characteristics of the head and shoulders pattern that are sought by traders. Edwards and Magee (2001, p.57) give a representative exposition. Four key characteristics are determined as necessary for a pattern to qualify. These pertain to the formation of the left shoulder, head, right shoulder and 'confirmation' by price crossing the 'neckline'. The neckline is defined by a line drawn to connect the troughs either side of the head and extended rightwards. These criteria are summarised in Table 3.1.

Crucially, Murphy (1999) makes the distinction between head and shoulders *reversal* and *continuation* patterns. We will concentrate on the reversal pattern which, as the name implies, contains information signalling the reversal of a pre-existing trend. This is the most popular use of the head and shoulders by technical analysts, but the distinction is an important one and ignored in previous studies. Murphy agrees with the characteristics outlined above, and emphasises the importance of the neckline to 'confirm' signals from patterns. Similar to the distinction between continuation and reversal patterns, existing work has ignored the neckline with the exception of Savin et al. (2007). However, as discussed previously, there are lim-

itations in their approach and methodology that are addressed by the innovations in this work. Most importantly, they do not investigate the head and shoulders bottom formation, which is accorded equal importance by technical analysts to the top formation.

Other practitioner texts also broadly agree with the important features of the head and shoulders discussed above, particularly the significance of the neckline and the distinction between continuation and reversal patterns. Stevens (2002) illustrates the key features of the head and shoulders pattern, and agrees with the importance of prior trend. He states that “Head and shoulders patterns, as is true of other top and bottom patterns such as double and triple tops or bottoms, are more likely to occur after a trend has been underway for some time” (pp. 165-66).

The prominent technical analyst and commentator Martin Pring also makes clear the importance of the neckline (Pring, 1985, 1998). However, he also makes an important link between the time that patterns take to form and their economic value. Specifically, he states that “the longer it takes to form the pattern, other things being equal, the greater its importance” (Pring, 1998, p. 64). Further evidence that provides additional support for the above analysis can be found in Bulkowski (2005) and Kaufman (2005).

3.4 Data and methodology

3.4.1 Nature and breadth of data

The previous chapter noted the advantages of using a large dataset of individual stocks to study the profitability of a head and shoulders trading strategy. Most importantly—and particularly relevant to this chapter—the practitioner literature often concentrates on patterns occurring in individual stocks. Furthermore, use of a large dataset enhances our ability to draw inference about the success or otherwise of the pattern in producing excess returns. Therefore, this chapter also uses the same large sample of UK stocks. Specifically, the sample runs from January 1 1980 to December 31 2003. It is also worth remembering that dead stocks are included to avoid survivorship bias.

3.4.2 Identifying peaks and troughs in price data

The head and shoulders pattern is derived from a series of localised peaks and troughs in price data. In the previous chapter, kernel regression was used to smooth the noisy price data. The signum function was then employed to extract turning points in the first derivative of the smoothed series. In doing so, an alternating series of peaks and troughs could then be used to identify the occurrence of head and shoulders patterns.

As Chapter 2 showed, kernel regression is theoretically and practically very well suited to extracting useful information from noisy price data. Therefore, it continues to be used in this chapter. Given the work in Chapter 2 showed that locally optimised bandwidth appeared to be a superior approach, we use this method here.

3.4.3 Identifying head and shoulders patterns

As in the previous chapter, once peaks and troughs—local maxima and minima—have been determined, the geometric ‘definition’ of the head and shoulders pattern can be fitted to them, and in doing so patterns can be identified. As before, great care is taken to ensure that look-ahead bias is avoided. Furthermore, one of the valuable contributions in this study is the creation of the new concept of the trade lag. This allows measurement of the time elapsed between the completion of a pattern and our ability to detect it. In doing so, it is possible to investigate if ‘fresher’ patterns perform better.

Due to the much improved specification of the geometric formations of head and shoulders patterns in this chapter, the steps to identify a pattern and record buy and sell trades can be extended. In the previous chapter, the specification for the head and shoulders pattern following that employed by Lo et al. (2000) was applied to a trading strategy. The rules for identifying a head and shoulders (HS) and inverse head and shoulders (IHS) were specified as

$$\begin{array}{l}
 \text{HS A} \left\{ \begin{array}{l}
 HS1 > HS2, HS3 > HS2, HS3 > HS4, HS5 > HS4 \\
 HS3 > HS1, HS3 > HS5 \\
 HS1 \text{ and } HS5 \text{ within } 1.5\% \text{ of their average} \\
 HS2 \text{ and } HS4 \text{ within } 1.5\% \text{ of their average}
 \end{array} \right. \\
 \\
 \text{IHS A} \left\{ \begin{array}{l}
 HS1 < HS2, HS3 < HS4, HS3 < HS4, HS5 < HS4 \\
 HS3 < HS1, HS3 < HS5 \\
 HS1 \text{ and } HS5 \text{ within } 1.5\% \text{ of their average} \\
 HS2 \text{ and } HS4 \text{ within } 1.5\% \text{ of their average}
 \end{array} \right.
 \end{array}$$

It can be seen that IHS A and HSB take into account the alternating pattern of peaks and troughs that characterise the head and shoulders pattern.⁴ In addition,

⁴For brevity, when later referring to the specifications for buys and sells together we abbreviate to (I)HSB.

a constraint that $HS3 > HS1$ and $HS3 > HS5$ is employed to capture the head for inverse head and shoulders patterns and head and shoulders patterns. For head and shoulders patterns this is $HS3 < HS1$ and $HS3 < HS5$. In addition, $HS1$ and $HS5$, as well as $HS2$ and $HS4$, should be within 1.5% of their average. As discussed previously in chapter 2, this constraint requires all patterns to have a degree of vertical symmetry.

The first specification tested in this chapter makes some pivotal changes to this specification. First, the review of the practitioner literature shows that a prior trend should be in force for a head and shoulders reversal pattern to be recognised. As detailed above, this is a major shortcoming of existing research because traders use head and shoulders and inverse head and shoulders reversal patterns to forecast a price reversal. Thus, the implied response to the occurrence of a head and shoulders pattern is to open a short trade, and a long trade for an inverse head and shoulders patterns. Conversely, continuation patterns forecast a continuation in the prevailing trend. In response to a continuation, head and shoulders pattern traders would therefore open or accumulate long (not short) positions. For an inverse head and shoulders pattern, a short (not long) position would be opened or increased.

To take account of this often ignored, but vital, distinction, we require a prior uptrend to be in place for a head and shoulders pattern to initiate a short sale. Otherwise, the pattern could be a continuation or reversal pattern, but the restriction means that we focus attention on the more important reversal patterns. Similarly, for an inverse head and shoulders pattern, a prior downtrend must be in place. Whether a prior uptrend or downtrend is in place is measured by investigating whether the price at the start of a rolling window is lesser or greater than the first peak or trough in the pattern, for head and shoulders tops and bottoms respectively.




Figure 3.1: Diagram of an artificial head and shoulders pattern

$$\text{HS B} \left\{ \begin{array}{l} HS1 > HS2, HS3 > HS2, HS3 > HS4, HS5 > HS4 \\ HS3 > HS1, HS3 > HS5 \\ HS1 \text{ and } HS5 \text{ within } 1.5\% \text{ of their average} \\ HS2 \text{ and } HS4 \text{ within } 1.5\% \text{ of their average} \\ \text{A prior uptrend should be in place} \end{array} \right.$$

The counterpart inverse head and shoulders (head and shoulders bottom) pattern can be characterised by

$$\text{IHS B} \left\{ \begin{array}{l} HS1 < HS2, HS3 < HS4, HS3 < HS4, HS5 < HS4 \\ HS3 < HS1, HS3 < HS5 \\ HS1 \text{ and } HS5 \text{ within } 1.5\% \text{ of their average} \\ HS1 \text{ and } HS5 \text{ within } 1.5\% \text{ of their average} \\ \text{A prior downtrend should be in place} \end{array} \right.$$

Whilst looking for a prior trend is vital in terms of separating continuation and reversal patterns, and is a novel development, we also look at another important feature of the head and shoulders pattern that is recognised by technical analysts: the neckline. This is seen as an important confirmatory signal. We therefore extend HSB to be

$$\text{HS C} \left\{ \begin{array}{l} HS1 > HS2, HS3 > HS2, HS3 > HS4, HS5 > HS4 \\ HS3 > HS1, HS3 > HS5 \\ HS1 \text{ and } HS5 \text{ within } 1.5\% \text{ of their average} \\ HS2 \text{ and } HS4 \text{ within } 1.5\% \text{ of their average} \\ \text{A prior uptrend should be in place} \\ \text{Price should 'break' the neckline} \end{array} \right.$$

The additional restriction states that price should break the neckline before a trade is entered. This is achieved by requiring that price at the end of the rolling window is below that of the trough between the head and right shoulder. This is

illustrated in Figure 3.1. The counterpart inverse head and shoulders (head and shoulders bottom) pattern can be characterised by

$$\text{IHS C} \left\{ \begin{array}{l} HS1 < HS2, HS3 < HS4, HS3 < HS4, HS5 < HS4 \\ HS3 < HS1, HS3 < HS5 \\ HS1 \text{ and } HS5 \text{ within } 1.5\% \text{ of their average} \\ HS1 \text{ and } HS5 \text{ within } 1.5\% \text{ of their average} \\ \text{A prior downtrend should be in place} \\ \text{Price should 'break' the neckline} \end{array} \right.$$

The returns and excess returns from head and shoulders and inverse head and shoulders patterns are computed in the same way as Chapter 2. In addition, the rolling window approach that was previously described is also employed here. However, given that the above analysis proposed that traders may use longer periods of time to look for head and shoulder patterns, and those that form over longer periods have greater 'significance', in addition to a window length of 35 days, we also investigate returns with a window length of 65 days.

3.4.4 Conclusions

The methodology of this chapter has been developed to reflect its additional contribution. Most importantly, a specification of head and shoulders patterns is developed to take into account two important features recognised by technical analysts. These are the neckline and separating reversal and continuation patterns. Furthermore, a longer formation period is investigated. The trade lag concept developed previously is also used to aid in the analysis and bootstrap testing is performed.

3.5 Empirical results

This section of the chapter reports and analyses the returns of head and shoulders technical trading strategies. Whilst the findings are presented in a similar tabular fashion to those in Chapter 2, the trading rules are quite different. One of the main aims of this chapter is to investigate the profitability of a head and shoulders patterns of the type actively employed by traders. To achieve this, the current section shows the results for the two specifications, B and C, that were derived above from careful study of the practitioner literature. In addition, the analysis is further developed by allowing trading patterns to form over a longer period. As excess returns are evaluated for holding periods of 1, 5, 10, 20, 30 and 60 days, this allows us to gauge whether patterns established over an increased time provide more economically valuable information. Furthermore, the new technique developed in Chapter 2, the trade lag, is employed to assess the speed at which useful information in patterns decays.

The results from a trading strategy based on pattern specification HSB and IHSB are presented in Table 3.2. This specification requires that a prior uptrend (downtrend) be in place for a head and shoulders (inverse head and shoulders) to be recognised. Without this restriction, it is unclear whether a head and shoulders continuation or reversal pattern has been found. This is an especially important addition, as the response to continuation and reversal patterns by traders is completely different. Specifically, a head and shoulders bottom continuation pattern means that a traders will sell short, yet a head and shoulders bottom reversal pattern means that a long position should be opened.

Given the additional restriction, to separate reversal from continuation patterns, it is first interesting to compare the number of patterns identified under (I)HSB as opposed to (I)HSA. Looking back to Table 2.6 and comparing the results shows that around half of the patterns previously identified as head and shoulders are filtered out. This is important as it casts doubt on the findings of previous studies that fail to

Table 3.2: Head & Shoulders returns for 1980-2003 under pattern specification HSB and IHSB.

Period	N		Mean π		$\pi > 0$		<i>t</i> -statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	15224	18524	-0.0013 (2.1807)	-0.0133 (2.0388)	0.45	0.67	-0.0719 (0.9427)	-0.8853 (0.3760)	0.5210 (0.6024)
5	15224	18524	-0.0159 (4.5717)	-0.0002 (4.6203)	0.51	0.61	-0.4288 (0.6680)	-0.0073 (0.9942)	-0.3112 (0.7557)
10	15224	18524	-0.0986 (6.4374)	0.0682 (6.3949)	0.51	0.59	-1.8855 (0.0594)	1.4506 (0.1469)	-2.3771 (0.0175)
20	15224	18524	-0.5592 (9.4818)	0.2080 (8.7209)	0.50	0.57	-7.2511 (0.0000)	3.2369 (0.0012)	-7.7241 (0.0000)
30	15224	18524	-1.0055 (11.7529)	0.3098 (10.4679)	0.49	0.55	-10.5076 (0.0000)	4.0099 (0.0001)	-10.8468 (0.0000)
60	15224	18524	-1.2695 (16.7620)	-0.0946 (15.5702)	0.51	0.54	-9.2475 (0.0000)	-0.8196 (0.4125)	-6.6593 (0.0000)

The sample period is January 1 1980 to December 31 2003 comprising the 350 largest stocks (resampled annually) by market capitalisation. Dead stocks are included. Smoothing of the price series to allow for identification of head and shoulders patterns is performed with kernel smoothing using a locally optimised bandwidth. Pattern specification HSB and IHSB introduce the restriction that a prior uptrend or downtrend should be in place, respectively. This allows reversal patterns to be distinguished from continuation patterns. "Period" is the holding period, i.e. 30 would represent the return from t_1 to t_{30} where t_1 is the buy date. "N Buy(Sell)" represents the number of buys(sells). "Mean π Buy (Sell)" is the mean return for buys(sells), with standard deviation shown below in parentheses. " $\pi > 0$ " shows the percentage of profitable trades for buys and sells (the 'hit rate'). *t*-statistics are shown with *p*-values below in parentheses. Buy-Sell reflects a standard *t*-ratio for the difference between mean buy and sell trade returns. For ease of reference, these results are reported as direct percentages.

make the distinction between continuation and reversal patterns. Accordingly, the results obtained are tested against incorrect hypotheses given that the continuation and reversal patterns forecast a price move in completely opposite directions.

However, the mean excess returns from patterns (I)HSB do not represent a profitable technical trading strategy. Without taking into account a prior uptrend or downtrend, although relatively small in magnitude, the mean sell returns were all negative across holding periods of 1 to 60 days. Mean buy returns were positive at 1 and 5 days and negative for longer trade times. By comparison, the mean buy returns presented in Table 3.2 are negative for all holding periods. Indeed, the mean buy return for the 60 day holding period is -1.3% . The mean sell returns provide little further support. Whilst the sell returns over 1, 5 and 60 days are negative, they are smaller in size than without the prior trend restriction. Furthermore, sell returns at 10, 20 and 30 days are positive. Although these results clearly do not form the basis of a profitable trading strategy, the t statistics show that five of the individual buy/sell mean returns are significant at the 5% level. The difference between buys and sells is also statistically significant for holding periods of 10, 20, 30 and 60 days. This suggests that the head and shoulders patterns are providing information, but the mean returns show that price does not move in the forecast direction.

Chapter 2 introduced the concept of the trade lag, which allowed measurement of how quickly profits from the trading strategies based on head and shoulders patterns decay. In application, this demonstrated that the most recent patterns produced far more profitable buy and sell trades, although the returns from the buys reversed after 10 days of holding the long position. Given the value of this approach, Table 3.3 imposes a trade lag of ≤ 5 days for patterns detected under (I)HSB. As before, a large reduction in the number of patterns is seen. However, there is a clear difference compared to the application in Chapter 2. The number of buy and sell trades, from inverse head and shoulders and head and shoulders

Table 3.3: Head & Shoulders returns for 1980-2003 under pattern specification HSB and IHSB with a trade lag of ≤ 5 .

Period	N		Mean π		$\pi > 0$		t-statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	3995	5442	0.0473 (1.8502)	-0.0765 (1.8763)	0.46	0.77	1.6151 (0.1064)	-3.0079 (0.0026)	3.1837 (0.0015)
5	3995	5442	0.1601 (4.4286)	-0.2691 (5.0314)	0.51	0.71	2.2810 (0.0226)	-3.9427 (0.0001)	4.3006 (0.0000)
10	3995	5442	-0.0198 (6.4123)	-0.3768 (7.0726)	0.50	0.70	-0.1945 (0.8458)	-3.9234 (0.0001)	2.5187 (0.0118)
20	3995	5442	-0.4234 (9.1639)	-0.2999 (9.2356)	0.49	0.68	-2.9057 (0.0037)	-2.3886 (0.0169)	-0.6442 (0.5195)
30	3995	5442	-1.0738 (11.5912)	-0.1250 (10.3472)	0.48	0.65	-5.8187 (0.0000)	-0.8862 (0.3755)	-4.1782 (0.0000)
60	3995	5442	-1.7960 (17.0940)	-0.2856 (15.3402)	0.49	0.61	-6.5615 (0.0000)	-1.3609 (0.1736)	-4.4975 (0.0000)

patterns respectively, is no longer similar. There are 3,995 buy trades and 5,442 sell trades. This may explain the differing pattern between the mean buy and sell returns. It is seen that, compared with the results where no trade lag is in place, the mean buy returns for 1 and 5 days have now become positive as predicted by the inverse head and shoulders pattern. Furthermore, the magnitude of the negative mean excess returns for the other time horizons has reduced. In terms of sell trades, all of the mean returns are now in the expected direction. For instance, the mean 20 day excess sell return with no trade lag is -2.6% per annum (i.e. negative excess return). With a trade lag of ≤ 5 days, the mean sell return is around +3.7% per annum. Trading only the 'freshest' patterns therefore again makes a clear difference to returns.

It is interesting to compare the magnitude of returns from (I)HSA and (I)HSB with a trade lag of ≤ 5 to see if the introduction of the prior uptrend/downtrend restriction to detect reversal (and not continuation) patterns has had an impact. Focussing only on strategies that are significant, the result is that imposing this restriction has not led to increased profitability in terms of buy or sell trades. This is

an interesting result because the (I)HSB specification, informed by the practitioner literature, should better capture the activities of traders. The head and shoulders pattern is essentially comprised on a sequence of peaks and troughs. This finding may suggest that the pattern is not performing in quite the way that traders think, and it may be that it is best at identifying areas of important support and resistance.

Table 3.4: Head & Shoulders returns for 1980-2003 under pattern specification HSB and IHSB with a formation period of 65 days.

Period	N		Mean π		$\pi > 0$		<i>t</i> -statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	28586	39565	-0.0482 (2.1459)	0.0113 (1.9162)	0.44	0.76	-3.7979 (0.0001)	1.1746 (0.2402)	-3.8033 (0.0001)
5	28586	39565	-0.2975 (4.9300)	0.0895 (4.2521)	0.49	0.68	-10.1991 (0.0000)	4.1834 (0.0000)	-10.9496 (0.0000)
10	28586	39565	-0.6863 (7.4276)	0.1620 (5.8996)	0.49	0.67	-15.6079 (0.0000)	5.4523 (0.0000)	-16.5624 (0.0000)
20	28586	39565	-1.2290 (10.6710)	0.1007 (8.6471)	0.49	0.65	-19.4222 (0.0000)	2.3090 (0.0210)	-17.8962 (0.0000)
30	28586	39565	-1.5122 (12.9881)	-0.0321 (10.8577)	0.49	0.64	-19.6047 (0.0000)	-0.5844 (0.5590)	-16.1301 (0.0000)
60	28586	39565	-1.6826 (17.6768)	-0.5501 (16.3964)	0.49	0.62	-15.9598 (0.0000)	-6.6062 (0.0000)	-8.6040 (0.0000)

One of the important outcomes from the study of the practitioner literature is that the length of time over which patterns form, as employed in previous work, may be too short. To investigate this, the (I)HSB pattern specification is also evaluated with a formation period of 65 days. Table 3.4 presents the results from this empirical work. It is immediately apparent that the number of patterns identified is over double the amount observed with a formation period of 35 days. The results show that this has not, however, resulted in a more profitable trading strategy. Again, all the mean buy excess returns are negative, but are larger in magnitude (i.e. losses have increased). For instance, a formation period of 65 days and holding period of 60 days is -1.6826% against 1.2695% for 35 days. A similar effect is seen in the mean sell returns but now, for the longer formation period, the

1 and 5 day mean sell returns are also positive (against the head and shoulders predicted direction of the trades). It is interesting to note that all the buy and sell returns are now significant at the 1% level, with the exception of the 1 day sell return. Further, there is a statistically significant difference between the mean buy and sell returns at the 1% level.

Table 3.5: Head & Shoulders returns for 1980-2003 under pattern specification HSB and IHSB with a formation period of 65 days and a trade lag of ≤ 5 days.

Period	N		Mean π		$\pi > 0$		<i>t</i> -statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	3113	4826	0.0322 (1.9398)	-0.0981 (1.5821)	0.45	0.90	0.9256 (0.3547)	-4.3065 (0.0000)	3.2711 (0.0011)
5	3113	4826	0.0228 (4.8572)	-0.2270 (4.5438)	0.51	0.82	0.2615 (0.7937)	-3.4669 (0.0005)	2.3263 (0.0200)
10	3113	4826	-0.3079 (7.1507)	-0.0380 (6.4113)	0.50	0.76	-2.4020 (0.0164)	-0.4107 (0.6813)	-1.7496 (0.0802)
20	3113	4826	-0.9760 (10.0245)	0.1542 (8.4999)	0.48	0.75	-5.4199 (0.0000)	1.2560 (0.2092)	-5.3761 (0.0000)
30	3113	4826	-1.6134 (12.9885)	0.0943 (9.8504)	0.49	0.73	-6.9063 (0.0000)	0.6614 (0.5084)	-6.6211 (0.0000)
60	3113	4826	-1.9702 (18.4671)	-0.2709 (15.5868)	0.50	0.68	-5.8959 (0.0000)	-1.1952 (0.2321)	-4.4012 (0.0000)

Whilst increasing the formation period for the chart patterns has increased the size of returns, these are not in the anticipated direction and thus this does not constitute a profitable trading strategy. As the trade lag has been shown to filter some of the lesser performing patterns, Table 3.5 shows pattern specification (I)HSB with a trade lag of five or fewer days from pattern completion to identification. Very similar results are seen compared to imposing the same trade lag with the 35 day formation period. The mean buy returns at 1 and 5 days are now in the correct (positive) direction, and are significant at 1%, with significant negative mean sell returns at 1 and 5 days. The case of the one day mean sell return is particularly interesting. Out of the 4,826 trades, 90% were profitable. The mean excess return of -0.0981% is equivalent to approximately -24.5% annually. Whilst

over an apparently small number of trades, there is still just over one sell trade per stock, per annum, on average. When coupled with a success rate of 90% this is a seemingly successful result.

Table 3.6: Head & Shoulders returns for 1980-2003 under pattern specification HSC and IHSC.

Period	N		Mean π		$\pi > 0$		<i>t</i> -statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	7588	8698	0.0618 (1.8955)	-0.0583 (2.3682)	0.47	0.64	2.8389 (0.0045)	-2.2968 (0.0217)	3.5368 (0.0004)
5	7588	8698	0.1095 (4.2802)	-0.0933 (5.2262)	0.52	0.57	2.2277 (0.0259)	-1.6648 (0.0960)	2.6846 (0.0073)
10	7588	8698	0.0045 (6.0519)	-0.0887 (6.9906)	0.51	0.56	0.0650 (0.9482)	-1.1830 (0.2368)	0.9032 (0.3664)
20	7588	8698	-0.7835 (9.2865)	0.0717 (8.9878)	0.50	0.53	-7.3286 (0.0000)	0.7427 (0.4577)	-5.9576 (0.0000)
30	7588	8698	-1.3229 (11.6377)	0.1057 (10.7956)	0.49	0.52	-9.8621 (0.0000)	0.9101 (0.3628)	-8.1071 (0.0000)
60	7588	8698	-1.3579 (16.9699)	-0.9208 (16.4088)	0.51	0.54	-6.8978 (0.0000)	-5.1987 (0.0000)	-1.6689 (0.0952)

Table 3.6 shows results with a further restriction; price now needs to have crossed the neckline for a trade to take place. Technical analysts place great importance on this. These results from (I)HSC are best compared with those from (I)HSB displayed in Table 3.2. Similar to the move from (I)HSA to (I)HSB, filtering patterns to only look at instances where the neckline is crossed has led to far fewer being identified. There are roughly half the number of patterns with this restriction. The significance attached to the neckline by traders seems partially justified. The returns from buys at 1 and 5 days are now in the expected direction, with excess returns of 0.0618% and 0.1095%, respectively. The latter corresponds to an excess return of around 5.5% annually. Similarly, for 1 and 5 days, the sell returns have increased in size, although only the 1 day return is significant. However, the picture is somewhat mixed. Trades held for 20-30 days still produce returns of the opposite sign to that predicted by the head and shoulders pattern. However, it

seems to be the case, particularly at 1 and 5 days, that the neckline is important, and the importance attached to it by technical analysts is warranted.

Table 3.7: Head & Shoulders returns for 1980-2003 under pattern specification HSC and IHSC with a trade lag of ≤ 5 and pattern formation period of 35 days.

Period	N		Mean π		$\pi > 0$		<i>t</i> -statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	1521	2110	0.1758 (1.6834)	-0.1447 (2.2929)	0.50	0.80	4.0728 (0.0000)	-2.8992 (0.0038)	4.6135 (0.0000)
5	1521	2110	0.4540 (4.2253)	-0.5060 (6.0551)	0.55	0.73	4.1852 (0.0000)	-3.8385 (0.0001)	5.2987 (0.0000)
10	1521	2110	0.1595 (6.0299)	-0.6637 (8.4496)	0.54	0.71	1.0300 (0.3032)	-3.6057 (0.0003)	3.2450 (0.0012)
20	1521	2110	-0.6774 (8.9758)	-0.3905 (9.9600)	0.51	0.68	-2.9338 (0.0034)	-1.7984 (0.0723)	-0.8922 (0.3723)
30	1521	2110	-1.2701 (11.1582)	-0.1798 (10.6028)	0.49	0.64	-4.4201 (0.0000)	-0.7762 (0.4377)	-2.9872 (0.0028)
60	1521	2110	-1.9312 (17.5134)	-0.7455 (15.6225)	0.50	0.64	-4.2521 (0.0000)	-2.1785 (0.0295)	-2.1436 (0.0321)

Table 3.7 follows the approach taken with (I)HSB and introduces a trade lag of ≤ 5 days. Imposing this filter leaves a very small number of trades over the sample period; just 1,521 buys and 2,110 sells are recorded. The change from introducing the trade lag is seen most clearly at the shortest holding period of 1 and 5 days. Both mean excess buy returns and mean excess sell returns are significant. The mean excess buy return over 5 days with the trade lag is 0.4540%, compared to 0.1095% without. The mean excess sell returns at the shortest time horizons are similarly increased in magnitude. These results show that the head and shoulders pattern, with account taken of the neckline, prior trend and looking at the most recently formed patterns constitutes a profitable trading strategy, at the shortest time horizons. However, this must be placed against the observation that only a relatively small number of trades per stock, per year, are recorded. Traders could probably not employ this strategy all the time.

As with (I)HSB, the neckline restriction is further evaluated by allowing patterns

Table 3.8: Head & Shoulders returns for 1980-2003 under pattern specification HSC and IHSC with a pattern formation period of 65 days.

Period	N		Mean π		$\pi > 0$		t-statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	14708	18157	0.0129 (1.8558)	-0.0105 (2.2372)	0.45	0.69	0.8430 (0.3993)	-0.6322 (0.5273)	1.0162 (0.3095)
5	14708	18157	-0.1689 (4.4317)	0.0149 (4.8238)	0.49	0.61	-4.6192 (0.0000)	0.4156 (0.6777)	-3.5599 (0.0004)
10	14708	18157	-0.5093 (6.7748)	0.0238 (6.6089)	0.50	0.60	-9.1046 (0.0000)	0.4851 (0.6276)	-7.1843 (0.0000)
20	14708	18157	-1.1391 (10.2903)	-0.0105 (9.2040)	0.50	0.57	-13.3825 (0.0000)	-0.1526 (0.8787)	-10.4652 (0.0000)
30	14708	18157	-1.3298 (12.5133)	-0.2275 (11.4548)	0.50	0.57	-12.8211 (0.0000)	-2.6653 (0.0077)	-8.3134 (0.0000)
60	14708	18157	-1.1771 (16.6054)	-1.3665 (17.6610)	0.50	0.59	-8.5105 (0.0000)	-10.3503 (0.0000)	0.9928 (0.3208)

to form over a longer time period of 65 days, as suggested by some of the practitioner literature. Table 3.8 shows (I)HSC with a pattern formation period of 65 days. As previously seen, according a greater amount of time for the occurrence of peaks and troughs to form patterns markedly increases the number of instances of patterns recorded. There are now roughly twice the number of patterns shown in Table 3.6. It is shown that the magnitude of returns is somewhat smaller than exhibited with a formation period of 35 days. Furthermore, fewer of the buy and sell excess returns exhibit a significant difference from zero. To provide a complete analysis, Table 3.9 shows the 65 day formation (I)HSC specification with a trade lag of ≤ 5 . As before, the number of pattern instances reduces considerably. Similarly, the magnitude of returns increases at short time horizons; the mean excess buy returns at 1 and 5 days are significant. Whilst these returns are large compared with previous results, this should not be accorded undue importance. Given that there were only 1,140 buys and 1,743 sells over the sample period, traders could not employ this strategy very often. This means that the seemingly large 1 and 5 day returns occur infrequently and would therefore almost certainly not result in a

profitable trading strategy overall.

Table 3.9: Head & Shoulders returns for 1980-2003 under pattern specification HSC and IHSC with a trade lag of ≤ 5 and a formation period of 65 days

Period	N		Mean π		$\pi > 0$		<i>t</i> -statistics		
	Buy	Sell	Buy	Sell	Buy	Sell	Buy	Sell	Buy-Sell
1	1140	1743	0.2118 (1.9295)	-0.1832 (1.7240)	0.49	0.89	3.7067 (0.0002)	-4.4362 (0.0000)	5.7042 (0.0000)
5	1140	1743	0.3178 (4.6993)	-0.4553 (5.1821)	0.56	0.81	2.2820 (0.0227)	-3.6682 (0.0003)	4.0507 (0.0001)
10	1140	1743	-0.0724 (6.2157)	-0.2176 (7.7649)	0.53	0.74	-0.3933 (0.6942)	-1.1698 (0.2422)	0.5300 (0.5962)
20	1140	1743	-1.0903 (9.3572)	-0.0512 (9.8465)	0.49	0.72	-3.9239 (0.0001)	-0.2163 (0.8288)	-2.8219 (0.0048)
30	1140	1743	-1.5061 (11.3015)	-0.2160 (10.4768)	0.49	0.71	-4.4759 (0.0000)	-0.8568 (0.3917)	-3.1283 (0.0018)
60	1140	1743	-2.1884 (17.8177)	-1.3164 (16.1313)	0.50	0.71	-4.0994 (0.0000)	-3.3834 (0.0007)	-1.3613 (0.1735)

3.5.1 Transaction and short selling costs

The previous chapter noted that transaction costs were important, and could negate the profitability of a head and shoulders trading strategy. The one-way break-even cost was discussed as a means of assessing this. Given that, as noted above, pattern specification (I)HSC does not appear to form the basis of a profitable trading strategy it is not necessary to consider transaction costs. However, (I)HSB shows significant mean excess returns for many holding periods.⁵

Table 3.3 shows positive mean excess buy returns at horizons of 1 and 5 days. The one-way break-even transaction costs are 0.0237% and 0.0801%, respectively. Clearly, this means that in both cases that excess returns net of transaction costs - even at the less conservative figure of 0.18% proposed by Jones (2002) - would be negative. All of the mean excess sell returns for (I)HSB with a trade lag of less than five days are negative (i.e. profitable). However, the break-even one-way

⁵As discussed in the previous chapter, it is also prudent to consider that short sale costs may have a bearing on profitability.

Table 3.10: Bootstrap results from 500 simulated series compared to the actual price series.

Holding Period	Fraction of simulations greater than actual series			
	Buy	Sell	σ_b	σ_s
1	0.0	0.0	100.0	100.0
5	0.0	0.8	0.0	0.0
10	100.0	0.6	0.0	0.0
20	99.6	1.0	0.0	0.0
30	99.2	0.4	100.0	100.0
60	99.0	1.2	99.0	100.0

As detailed in the methodology section, the original price series is ‘shuffled’ (resampled with replacement) 500 times. The algorithm for identifying head and shoulders patterns is run on these pseudo price-series. For the mean, buy, sell and standard deviations of buy and sell returns (σ_b and σ_s , respectively), the columns report the fraction of simulations greater than the original series. Results are presented for 1 to 60 day holding periods.

transaction cost (half of the mean excess return) is less than the more conservative 0.23% figure in all cases. At an assumed transaction cost of 0.18%, only the 10-day sell trade remains profitable.

3.5.2 Bootstrap tests

As in Chapter 2, bootstrap testing is employed to compare results from the original price series to 500 simulated series. These were constructed as a random walk. (I)HSB with a 35 day holding period and trade lag of ≤ 5 was selected as the best candidate for bootstrap analysis; the results for this specification were shown in Table 3.3. To allow an insight into the sufficiency of 500 simulations, Figure 3.2 gives an example of the convergence of the estimates as the number of replications is increased.⁶ As with Brock et al. (1992), these results show that the simulated

⁶When looking at this chart, recall that standard deviation is based on a multiple of $100 \times \log$ cumulative return.




Figure 3.2: Cumulative mean standard deviations over 500 simulations.

series 'settle' relatively quickly, giving support to the choice of 500 replications as a good balance between the extensive computational time needed to compile them and wanting to increase power.

Table 3.10 gives the findings of the bootstrap analysis. Interestingly, the results are almost completely consistent with those found for (I)HSA, that were presented earlier in Table 2.9. Recall that we can think of the values in the table as simulated p -values. Looking between the results from the bootstrapping of pseudo-series and the 'real' series shows that again the 1 and 5 day buy trades are significant based

on the bootstrapping. This conflicts with the standard t -statistics and suggests that the distributional assumptions for this test may be problematic. The buy returns for 20, 30 and 60 days are shown to be insignificant, which also disagrees with the t -statistics.

Including a filter to isolate reversal from continuation patterns now means that the simulated series provide support for the sell returns at all time horizons. It was noted previously that bootstrapping the standard deviation for buys and sells (σ_b and σ_s) gives insight into the riskiness of the trades signalled by the head and shoulders pattern. The results show that the buy standard deviations from none of the simulated series were greater than the original series at 5, 10 and 20 days. Volatility therefore does not appear to be a factor in explaining excess returns at these time horizons. However, we should also note that the 1, 30 and 60 day holding period returns are subject to greater volatility than almost all the simulated series. The same picture is seen for the sell standard deviations.

The results from this bootstrap testing are interesting. They show the head and shoulders pattern does appear to be able to provide useful economic information and, crucially, at intermediate time horizons, the returns do not seem to be accompanied by commensurately higher volatility. However, as noted in Chapter 2, it would be beneficial for future research to further examine this issue by looking at risk-adjusted returns using an appropriate factor model.

3.6 Conclusions

This chapter has argued that the head and shoulders patterns evaluated in previous work have not been sufficiently similar to those that would be identified by traders. This is important because the practitioner literature is specific about the features of patterns that technical analysts would require before trading. Accordingly, two new specifications, (I)HSB and (I)HSC, are developed. The former takes into account the prior price trend before the pattern. Head and shoulders patterns can be split clearly into two types—the reversal and continuation pattern—but this distinction has not previously been made. The latter also includes the neckline, which traders see as an important confirmatory factor, needed before opening a position. In addition, it is argued that the formation time of 35 days that has been tested in early empirical work is too short. Therefore, profitability of the head and shoulders pattern for formation periods of 35 and 65 days were investigated.

The empirical work also builds upon and considerably extends the developments made in Chapter 2. To this end, patterns are evaluated over a range of holding periods from 1 to 65 days. The trade lag is also employed to investigate how quickly profits from head and shoulders patterns decay. The findings presented in this chapter show that, to some extent, the head and shoulders pattern appears to form the basis of a profitable trading strategy.

The first group of findings related to the inclusion of a filter to separate reversal and continuation chart patterns. This led to about 50% less patterns being identified. The interpretation is that around half of head and shoulders formations signal reversal, and half continuation of the existing trend. This important distinction, which is recognised and given emphasis by traders, is not made in previous studies. With a trade lag of ≤ 5 days, the head and shoulders produced useful information, but only at the shorter holding periods. This suggests that profits from the trading pattern decay relatively quickly, and that the gap between being able to identify a pattern (when the relevant peaks and troughs form) is

important. However, the most interesting result is that restricting our focus to only continuation patterns does not improve the size of profits. This finding suggests that technical analysts may be unwise in attempting to classify head and shoulders patterns into continuation and reversal patterns.

The empirical analysis also looked at the profitability of head and shoulders patterns when a break of the neckline occurs. This is important because traders view this as an important 'confirmation' signal. The neckline seems to have a measurable impact on returns. For example, buying on the basis of inverse head and shoulders patterns and selling after 5 days produces an excess return in the order of 5.5% annually. However, returns seem to reverse relatively quickly, so that holding for longer periods above 20 days is not profitable.⁷

As well as investigating patterns for a formation period (and rolling window) of 35 days, results were also presented for a formation period of 65 days. This resulted in a large increase in the number of patterns identified. In the case of separating reversal and continuation patterns, this did not improve profitability. Indeed, mean excess buy and sell returns became smaller or negative. It is also the case that this longer formation period does not contribute to the success of trades undertaken when the neckline is broken.

Bootstrapping was conducted to evaluate the most interesting and important results. The results of this showed that the head and shoulders pattern provided significant returns; however, buys and sells performed differently. The inverse head and shoulders, for buy signals, produced significant returns for holding periods up to 5 days. By comparison, sell returns from the head and shoulders top provided significant returns up to 60 days. Most importantly, bootstrapping showed the head and shoulders patterns picked out lower standard deviations from the actual series than in all of the 500 bootstrap cases. This suggests that the

⁷The neckline filter was also investigated with a trade lag of ≤ 5 days but, although results supported the head and shoulders in terms of profitability, a relatively small number of trades were identified.

excess returns are not unduly risky.

Overall, these results show that the faith placed in the head and shoulders pattern by technical analysts is not misplaced. However, some of their assumptions should be questioned. In particular, attempting to identify continuation and reversal patterns. The findings of this study suggest that it is far more important to minimise the time between the identification of patterns and opening trades. However, traders are right to look for confirmation from a break of the neckline before entering trades. Whilst these results have implications for professionals, they are equally valuable in terms of market efficiency. We would not expect the head and shoulders pattern to be profitable as it employs nothing more than past price history. The bootstrap results for standard deviations suggest that, for the specifications of the head and shoulders showing the greatest excess returns, the trades identified are not accompanied by increased volatility.

Chapter 4

Intraday Reversal or Relative Strength?

4.1 Introduction

Following the seminal paper by De Bondt and Thaler (1985), a substantial amount of empirical work has demonstrated that investors tend towards overconfidence in their beliefs and, as a result, financial assets returns tend to display a ‘reversal’ effect. Thus a contrarian strategy of buying stocks that have previously performed poorly and selling stocks that have performed well is profitable. Conversely, Jegadeesh and Titman (1993) initiated what has become a very extensive literature on momentum in asset returns.¹ Momentum is based on the core idea of relative strength; relative strength strategies seek to buy stocks that have performed well in the past and sell stocks that have performed poorly, thus implying that a previous trend will continue. This is an important aspect of technical analysis, which is largely built on the concept of trends.

The time horizon over which portfolios are formed and their performance evaluated is critical in making sense of these seemingly conflicting schools of thought. First, reversal effects dominate from week to week and month to month (Lehmann, 1990; Jegadeesh, 1990; Antoniou et al., 2006). Second, relative strength (momentum) is prevalent at the intermediate term from three to 12 months (Jegadeesh and Titman, 1993, 2001). Third, reversal again dominates over the longer term of 3-5 years (De Bondt and Thaler, 1985, 1987).

Yet, despite voluminous research on reversal and momentum over periods of up to five years, we know very little about whether it is momentum or reversal—or neither—which characterises financial asset returns at intraday time horizons. This is somewhat surprising given the key role that time horizon has played in discerning between momentum and reversal effects.

However, the pervasiveness of these effects, which do not appear to be conditional on the asset pricing model selected or a particular sample period, pose a

¹Although Jegadeesh and Titman (1993) introduced relative strength to the modern finance literature, research on this subject has a long pedigree (for example Levy, 1967).

significant problem with respect to the efficient markets hypothesis.

One important factor motivating the study of price momentum is evidence suggesting that institutional investors and mutual funds buy past winners and sell past losers. Jegadeesh et al. (2004) show that analysts tend to favour recommending positive momentum stocks. In addition, Carhart (1997) identifies that outperforming mutual funds tend to hold stocks which have exhibited price momentum over the previous year. Yet, this result is not due to fund managers following a momentum strategy, but rather “because some funds just happen by chance to hold relatively larger positions in last year’s winning stocks” (p. 58). Furthermore, there is evidence to suggest that large investors tend towards following price momentum at higher frequencies. Using proprietary NASDAQ trade data, Griffin et al. (2003a) show that institutional investors tend to be intraday momentum traders. However, this result does not accord with Nofsinger and Sias (1999), who find that institutions tend not to be momentum investors.

A different picture emerges with respect to individual investors. Griffin et al. (2003a) show that, in contrast to institutions, individuals tend to be best characterised as contrarians. Overall, there is some evidence to suggest that overreaction may be present on an intraday basis. For example, Fabozzi et al. (1995) use a filter to isolate large price changes in intraday data for NYSE and AMEX listed stocks in 1989. A tendency for large price changes to be reversed during the trading day is shown. Grant et al. (2005) find a similar reversal effect to be prevalent in the S&P 500 futures. However, the existing research does not set out to comprehensively examine intraday momentum and reversal effects. So, we lack evidence about the profitability of intraday momentum and reversal strategies.

Providing further insight into this disagreement, in a study using trade and quote data, Hvidkjaer (2006) investigates not whether a momentum effect exists at short time horizons, but how investors trade in momentum portfolios. Results show that momentum may be driven by underreaction amongst small traders, an

effect not present among large traders.

The connection between momentum and institutions, funds and analysts provides further strong motivation to investigate momentum using previously untapped ultra high-frequency data.

A possible mediator between indicators of intraday momentum and reversal seen over time may be the rapidly increasing trading volume on major world markets, and the related phenomenon of day trading. The turnover rate (shares traded as a percentage of shares outstanding) for the NYSE was 99 per cent in 2003 as opposed to 54 per cent in 1993.²

This may help to bridge the gap between evidence, on the one hand, of some institutional investors being driven by momentum, and, on the other, the limited evidence of intraday overreaction. Due to their remit, most mutual funds do not seek to capture intraday profits, yet this is the exact goal of day traders. There is evidence of overreaction generally by investors, for example Odean (1998) and Barber and Odean (1999). Specifically focussing on investors with short time horizons, Mizrach and Weerts (2009) capture 'chatter' between day traders on an internet forum and analyse the results to determine what drives such traders, what strategies they adopt and how profitable they are. They note that the majority of day traders use momentum to place trades; indeed, "the survey showed that 75 percent pick a stock and its entry point based on momentum measures" (p. 269). This evidence does not, on the face of things, seem to accord with many investors succumbing to overconfidence. Given this lack of clarity, it is important to establish whether momentum or reversal is profitable intraday. In doing so, it is possible to gain an insight into whether day traders, such as those evaluated by Mizrach and Weerts, are misguided in their strategies.

Given the increasing prominence of short term trading, as evidenced by the increase in turnover and program trading, it is important to gain knowledge of

²As reported by the NYSE at <http://www.nyxdata.com>.

intraday momentum and reversal behaviour. The core objectives of this chapter are to, first, provide the first comprehensive study of high-frequency momentum and reversal strategies, and evaluate their profitability. Second, to investigate how the profitability of these strategies interacts with the month of the year, day of the week and hour of the trading day. Third, evaluate portfolios formed on size, to examine if profitability is conditional upon this factor.

Contributing to our understanding of reversal and relative strength in a completely new way, this chapter investigates short-term relative strength and contrarian strategies with ultra high-frequency data. Using trade data from the NYSE trade-and-quote (TAQ) database for the constituents of the S&P 500 from January to December 2005, it is found that buying stocks with (relatively) low returns and selling stocks with (relatively) high returns over the previous 10 to 60 minutes forms a profitable trading strategy. Conversely, there is no evidence of price momentum—buying winners and selling losers—being profitable at short-term intraday time horizons. Results from this strategy are analysed based on month, day of the week, time of day and market capitalisation quintiles. In all cases, it is found that a reversal effect prevails.

Using high-frequency trade data from the New York Stock Exchange time-and-quote (NYSE TAQ) database, this chapter provides a completely new perspective on relative strength by investigating its intraday profitability. Whilst existing research is generally polarised between testing for momentum or reversal, this chapter takes a holistic approach and examines evidence for both effects. With a large sample of S&P 500 constituent stocks for the whole of 2005, the returns from a momentum strategy of buying stocks that have performed well in the previous 10-60 minutes and selling stocks that have performed poorly in the same period are established. Analysis clearly shows that a contrarian (reversal) strategy—rather than momentum—prevails. This has important implications for the large number of investors that trade intraday.

4.2 Literature

This chapter makes a distinctive contribution by investigating the intraday profitability of momentum and reversal strategies. Whilst Jegadeesh and Titman (1993) inducted momentum into the literature, there have been a large number of subsequent publications that document further developments in terms of developing the momentum strategy itself, investigating sources of profitability and extending analysis to markets outside the United States.

Much research in finance has been concerned with the attempt to establish whether financial asset returns are negatively or positively correlated with previous returns. One body of empirical work has demonstrated the existence of contrarian profits, accruing from buying losers and selling winners. On the other hand, there is also convincing evidence that buying winners and selling losers produces abnormal returns. Time horizon is important in bridging the gap between such reversal and relative strength (or momentum) effects. This chapter take an agnostic approach based on a starting point of evaluating whether relative strength or reversal effects prevail at short time horizons. This is also salient because, as noted above, there is evidence of differences in behaviour between groups of traders. Accordingly, it is important to evaluate the literature relating to both aspects. Since reversal and overreaction were generally investigated and documented first, this is chosen as the starting point for this review.

There is much evidence that the returns of individual stocks reverse over longer time periods, measured in years. De Bondt and Thaler (1985) studied portfolios comprised of long-term winners and losers. Using NYSE data from 1926-1982, they formed portfolios based on the performance of stocks over one to five years. The core result is that loser portfolios markedly outperform the market, whereas the winner portfolios underperform. Therefore a contrarian strategy can be profitably employed to exploit the tendency of winners (losers) to 'reverse' their gains (losses) in the future.

Several important aspects of these results have since become regarded as stylised facts. First, the post-formation performance of winners and losers is not symmetrical; indeed, the “overreaction effect” is substantially larger for losers than winners. Second, most of the excess returns accruing from the reversal strategy owe to especially good performance in January. Equally important in reconciling these results with the later momentum literature is that “the overreaction phenomenon mostly occurs during the second and third year of the test period” (De Bondt and Thaler, 1985, p.799). In other words, most of the price reversal effect occurs after the first year.

De Bondt and Thaler (1987) investigate a number of these unresolved issues; perhaps most importantly, the possible role of size and risk factors in explaining reversal patterns is addressed. These two explanations are not supported by the data. Furthermore, the general hypothesis of overreaction is strengthened by the pattern of earnings reported by winner and loser firms. This leaves the important conclusion that investors systematically overreact, placing more weight than is rational on the most recent information when making investment decisions.

Jegadeesh (1990) provides evidence questioning whether overreaction is important, irrespective of time horizon. His methodology is different from that of De Bondt and Thaler (1985) in that he studies monthly returns (rather than lagged yearly returns). However, this does not affect the key result: that there is highly significant negative autocorrelation in stock returns from one month to the next. However, there is also significant *positive* autocorrelation over longer timeframes, up to 12 months. January still appeared to be important, but not essential, in driving these results. This lends support to the time horizon sensitive nature of reversal and relative strength noted by De Bondt and Thaler (1985, 1987).

Further shortening the time horizon, Lehmann (1990) looks at stock return data from one to 52 weeks. The finding of reversal effects from one month to the next found by Jegadeesh (1990) is also found here. Controlling for the bid-ask spread

provides an element of robustness to these conclusions.

Jegadeesh and Titman (1993) shift the focus back from reversal to continuation.³ This study occupies the gap between evidence of long-term reversal and short-term reversal. A motivating factor for looking at the movements of winner and loser portfolios, formed over three to 12 months, is that this is deemed to tally with the length over which market participants employing relative strength make decisions.⁴ Jegadeesh and Titman use a formation period of between three and 12 months to construct decile portfolios, and evaluate performance over the ensuing three to 12 months. The results present strong evidence of price momentum; for example, a six-month symmetrical formation/holding period strategy gives an average compound excess return of just over 12 per cent annually.

However, these results do not directly contradict evidence of price reversal. In fact, the opposite is the case. The cumulative returns from a relative strength trading strategy decay rapidly after 12 months. Furthermore, there is evidence of a reversal effect in the first month following portfolio formation. Therefore, it would seem that stock returns display both relative strength and reversal effects depending on the time period over which one studies them.

A large body of research has expanded to assess momentum profits in differing time periods and in different markets. Rouwenhorst (1998) looks at twelve European markets, following the methodology of Jegadeesh and Titman, and finds that winners outperform losers with a return of approximately 1 per cent per month. Whilst these are developed markets, Rouwenhorst (1999) confirms a momentum strategy is profitable in emerging markets. In contrast, Hameed and Kusnadi (2002) find that for six Asian markets, there is no evidence to support a successful momentum trading strategy. Chan et al. (2000) survey the indices

³I say 'back' because, as Jegadeesh and Titman (1993) themselves point out, very early studies showed the success of relative strength strategies (for example Levy, 1967; Jensen and Benington, 1970).

⁴However, this is likely not be as true today given the increasing accessibility of markets, lower transactions costs and day traders.

of 23 markets for momentum effects, finding that profits are economically and statistically significant. This study does not construct momentum portfolios *per se*, but buys/sells based upon past performance relative to the other indices in the study; as such, it is a strategy closer to the original definition of relative strength. Griffin et al. (2003b) provide a wide-ranging study of momentum in international markets, showing clear stability of profits over time that do not appear to be conditional on the stage in the business cycle. Jegadeesh and Titman originally looked at NYSE and AMEX stocks from 1965-1989. Jegadeesh and Titman (2001) update the data and show a continuation of momentum profits into the 1990s. This is a strong response to any criticism of data mining in the original results. Widespread confirmation of momentum profits therefore exists, but no investigation has taken place on intraday data.

There have been a large number of extensions to momentum, linking the anomaly with other areas of finance. Lee and Swaminathan (2000) add trading volume, finding that adding volume as a criteria in portfolio formation improves momentum profits. Moskowitz and Grinblatt (1999) study momentum in relation to industry grouping. The study shows that whilst profits from pursuing a momentum strategy appear small, once industry is controlled for, there are larger and significant returns from buying (selling) winning (losing) industry groups. Martin and Grundy (2001) disagree that industry effects are a primary cause of momentum profits. They demonstrate the stability of momentum profits over a long time period from 1926, determining that factor models are very successful in explaining variability in winner and loser returns but not their mean returns. Relating momentum to analyst coverage, Hong et al. (2000) show that profits from a momentum strategy are greatest amongst firms with low analyst coverage, and profitability is dependent on firm size. Liew and Vassalou (2000) investigate whether momentum can be used as a predictor for economic growth, although the

results do not indicate that this is the case.⁵ Chan et al. (1999) perform a two-way analysis, linking momentum strategies to earnings surprises and upward forecast revisions. It is found that combining these two elements is highly advantageous with significant profits over a 6-12 month time horizon. Such extensions of the original momentum proposition serve to demonstrate the continued importance of momentum. Whilst these issues are clearly interesting, and contribute to our understanding of momentum, they are beyond the scope of the current study. Indeed, issues such as predicting economic growth and the relation with analyst coverage are not applicable at the intraday time horizon.

Various explanations for the persistent success of momentum strategies have been proposed. Chan et al. (1996) attribute the success of medium-term momentum strategies to investor underreaction in response to earnings announcements. Hong and Stein (1999) propose a model where information diffuses gradually, meaning that underreaction is inherently present, which momentum traders can profitably exploit. Griffin et al. (2003b) investigate if macroeconomic risk can explain momentum profits in international markets, although they conclude that this is not the case.

In addition, behavioural explanations of momentum abound. Barberis et al. (1998) develop a framework that explains momentum in the context of investor over and under-reaction. Similarly, Daniel et al. (1998) relate the success of momentum (and other) anomalies to investors' tendencies towards overconfidence. Conrad and Kaul (1998) demonstrate the primacy of momentum at medium-term time horizons and cite cross-sectional variation in the mean returns of stocks as a driving factor of this result.⁶ Jegadeesh and Titman (2001) identify evidence of delayed overreaction as a major cause of momentum profits. Using nine years of additional out-of-sample data over the original study, the profitability of a momentum trading

⁵They also investigate book-to-market and size, finding that these factors are useful in predicting future GDP.

⁶Conrad and Kaul also show a contrarian strategy is profitable over a longer time horizon, albeit for only part of their sample period.

strategy is found to be sustained at around 1% per month.

There has, however, been no investigation of a fully intraday momentum portfolio trading strategy, although a limited amount of research employs intraday prices to different ends. Hvidkjaer (2006) uses intraday data with momentum strategies, but only to test trade imbalances on momentum portfolios formed over conventional (longer) time periods. Chakrabarty and Trzcinka (2006) use the NYSE TAQ database to determine if the momentum strategy of Jegadeesh and Titman is robust to different stock price databases. The results show that this is not the case because of the different ways that the TAQ and CRSP database handle delisting firms.

Without working within the framework of forming portfolios based on prior price momentum, several studies have documented intraday price reversals in opposition to a momentum effect. However, this research generally focuses solely on stock price indices. Grant et al. (2005) find evidence of intraday price reversals, but the survey is limited to the S&P 500 futures. This study uses filters based on the opening price gap of $\pm 0.10\%$, $\pm 0.20\%$ and $\pm 0.30\%$. Their relative strength strategy does not form portfolios. Instead, it is conducted within the framework of an event study. Cumulative abnormal returns are assessed, conditioned on the occurrence of an opening price gap equal or larger than the three filter sizes. In looking at intraday reversals of “large” opening price changes in the S&P 500 futures, we do not gain a broad understanding of intraday momentum and reversal effects. By contrast, forming portfolios based on the established momentum portfolio methodology for 500 of the largest US stocks gives us a much greater insight into short-term momentum and reversal.

Yu et al. (2005) conduct a study on intraday reversal and momentum effects using NASDAQ-100 futures index data. The methodology fits a multiple regression model and relates intraday returns to the previous day's intraday and overnight returns, and also looks at the effects of the previous day's and overnight returns

conditioned on the market state (bull or bear market) and day of the week. Empirical results show that both momentum and reversal effects can be identified, and the sign of the previous day's return and overnight return is important in this. However, this work does not give us an insight into momentum and reversal effects present for individual securities. Furthermore, the analysis is based on sixteen intraday periods (with endpoints 15 minutes apart) in contrast to the minute-by-minute data used in this study. Together with the formation of portfolios, we therefore seek to obtain a greater understanding in this work.

Fabozzi et al. (1995) provide evidence that stock prices exhibit intraday reversals following large price changes. Similar to Grant et al. (2005), an event study framework is adopted. A 2% filter rule is used to identify 'large' price changes amongst NYSE and AMEX stocks in 1989 with the result that reversals tend to occur very soon after the change but subsequently level out. In a similar vein, Fung et al. (2000) document intraday price reversals in the S&P 500 and HSI futures market. Large price changes at the open tend to subsequently reverse intraday, albeit this effect is more pronounced in the HSI than S&P 500 futures. The observations above are also applicable here; these studies provide valuable and interesting results, but a clearer understanding of momentum and reversal effects with high-frequency data on a large sample of stocks is needed.

Coming at the issue from a momentum rather than reversal standpoint, Lam et al. (2007) propose a study of intraday momentum; however, in this case, intraday momentum is based upon the difference between the opening and closing price of a stock for a trading day. The study is also limited by its confinement to 13 stock indices and not looking at individual securities. Significantly, by not looking at intraday data, this work does not address the question of intraday momentum profitability in a fashion that allows us to draw significant inference. By contrast, the current study adopts a large sample of stocks with high-frequency data for a complete year and fully constructs momentum portfolios.

A recent study of momentum using weekly data is provided by Gutierrez and Kelley (2008). This work supports conclusions reached much earlier (for example Lehmann, 1990) in also documenting clear reversals in weekly returns. However, instead of viewing weekly returns in isolation, the study looks at what happens to winner and loser portfolios from one to 52 weeks after formation. Using data from 1983-2003, it is found that the one-week reversal effect is completely subsumed by momentum when the same portfolios are evaluated after one year.

Figelman (2007) looks simultaneously at momentum and reversal effects over short, intermediate and long time horizons. For the short term, results are similar in nature to Jegadeesh (1990) and Lehmann (1990): over one month it is clear that reversal, not momentum, dominates. Again, for both intermediate-term and long-term formation and holding periods, the results are similar to previous studies. Over 12-months, momentum dominates but reversal effects are clear over 48-months. What is particularly instructive is that momentum and reversal effects are still present a long time after they were initially discovered.

Limited evidence of intraday momentum and reversal effects is provided by Schulmeister (2008). Whilst this study does not look at individual stocks, it does use 30-minute data for the S&P 500 spot and futures market. What is particularly interesting is that Schulmeister finds that the intraday profitability has been remarkably persistent over time. However, the methodology adopted looks to employ buy and sell trading signals based on relative strength. This means that it is not possible to directly infer whether a momentum or reversal effect predominates over shorter time horizons. Furthermore, we have no knowledge of the behaviour of individual security returns.

Overall, we remain unclear about the nature of intraday reversal and momentum. Existing research provides limited evidence that there may be both intraday reversal and intraday momentum effects. However, much work considers only a single stock index future or a limited number of stock index futures, does not view

the problem within the established momentum portfolio framework established by Jegadeesh and Titman (1993), uses (relatively) infrequent intraday data sampling, and/or only examines the intraday reversal of “large” price changes at the open in stock index futures. This research seeks to address all of these concerns. In constructing portfolios with high-frequency data for 500 large US stocks in the S&P 500, it is possible to gain a comprehensive insight into intraday momentum and reversal effects. Furthermore, as we know that many anomalies seem to exhibit patterns conditional on the month of the year and day of the week, the results of this study are broken down by month and day to investigate this possibility.

4.3 Data and methodology

NYSE trade and quote (TAQ) data was used from January 1 to December 31, 2005. The constituents of the S&P 500 index were identified monthly (to take account of additions and deletions from the index), and real-time trade data obtained for these stocks from the consolidated trades database. Observations were collected from 9:30 a.m. EST to 4:00 p.m. EST. The trades data were filtered so that 5-minute data was obtained, with the closest trade to each minute throughout the trading day being taken.⁷ A 5-minute interval was used because this provides a good compromise between ultra high-frequency data and a sample that is able to be computationally evaluated in a reasonable amount of time. The 5-minute return for an intraday period, d , can be defined as

$$r_{i,t,d} = \ln P_{i,t,d} - \ln P_{i,t,d-1} \quad (4.1)$$

where $R_{i,t}$ is the return for an intraday 5-minute period for a particular stock, i , on trading day t . $P_{i,t,d-1}$ is the price of the stock 5-minutes prior to $P_{i,t,d}$.⁸ A full

⁷In following this procedure, the resulting data set comprised 9,732,007 observations. Data was downloaded from the NYSE TAQ database via the WRDS service.

⁸Intraday 5-minute returns are similarly defined by Hol and Koopman (2002) and Marshall et al.

trading day therefore consists of 78 intraday returns.

Due to the ultra high-frequency data that is used in this study, it is important to ensure that it is filtered to remove bad records from the trade data. Leaving potentially erroneous data in the sample risks bias in the results and may also alter the properties of the series (such as autocorrelation). Brownlees and Gallo (2006) highlight some of the potential problems of not using cleaned TAQ data for analysis. Several important issues can be identified: the data as downloaded from the TAQ database may be mis-ordered; trades may be time-stamped outside trading hours; trades may be reported much later than they actually occurred, and there may be data recording errors. Mis-ordered trades may be corrected by sorting the data by time stamp, and this is carried out in this study. Trades reported outside market hours are removed from the sample with only trades between 9:30 a.m. EST to 4:00 p.m. EST being considered.

Trades that are reported to the tape later than they actually occurred (with other trades reported in the intervening period) are denoted in the NYSE TAQ database by the Sale Condition field (COND) taking a value of 'Z'. Such trades are removed from the sample here. Similarly, trades occurring in sequence but reported later (COND field 'O') are also eliminated. The TAQ data also possess a Correction Indicator field (CORR) which identifies later corrections to TAQ trade records. The sample in this study only adopts trades where CORR equals 0 or 1. Trades where CORR=0 are regular trades not subsequently corrected and CORR=1 are original trades subsequently corrected. In the latter case, the record of the trade is logged at the original time with corrected trade data. Dropping trades where CORR is not equal to 0 or 1 removes the relatively small number of observations that have been noted as erroneous by the NYSE, or where the trade was cancelled. Removing erroneous trades and the use of 5-minute sampled data ameliorate the problems inherent in using TAQ database data in this study. Upon the completion

(2008b).

of filtering, the data set comprised 9,684,840 observations.

In order to investigate intraday momentum and reversal effects, portfolios were constructed according to the methodology of Jegadeesh and Titman (1993). Following their notation, these were formed on the basis of the continuously compounded return over the previous J minutes, and held for K minutes. More formally, the total formation period return for an individual stock upon which portfolios are formed can be defined as

$$TR_{t,t-J} = \sum_t^{t-J} R_{i,t} \quad (4.2)$$

Portfolio formation, J , and holding period, K , time-horizons of 10, 20, 30 and 60 minutes were used. At the start of each period all stocks in the S&P 500 were ranked in ascending order on the basis of their J minute returns and then allocated to ten equally weighted portfolios on this basis. The top portfolio, therefore, contains stocks which have performed the best in the previous J minutes (winners), and the bottom portfolio those that have performed the worst (losers). Each portfolio is then held for K minutes, and the returns to holding each portfolio calculated, as well as the reward for this strategy: the return for the top minus the bottom portfolio.

$TR_{i,t}$, the total return on a stock purchased or sold-short through this strategy in a particular period, is

$$TR_{t,t+K} = \sum_t^{t+K} R_{i,t} \quad (4.3)$$

Positive returns imply momentum and negative returns imply reversal. Jegadeesh and Titman increase the power of their statistical tests by using overlapping portfolios. This study follows this approach. Accordingly, in any given time period, t , portfolios are held that are constructed at time t as well as in the previous

$K - 1$ periods.⁹

Table 4.1: **Descriptive statistics**

Summary statistics for the 5-minute returns series of S&P constituents in 2005 (adjusted monthly for additions and deletions to the index). * indicates significance at the 5% level, ** indicates significance at 1%.

N	9,684,840
Mean	-0.00000587
Std. Dev.	0.0017256
Skewness	-0.5420983**
Kurtosis	183.452556**
D -stat	0.1002**

Table 4.1 presents descriptive statistics for the 5-minute filtered data. The mean, standard deviation, skewness, kurtosis and Kolmogorov-Smirnoff (D -Stat) statistics are shown. The D -Stat test rejects normality at the 1% level; given this, it is not surprising that skewness and kurtosis are present in the intraday returns. The high kurtosis values are suggestive of occasional extreme movements between 5-minute return periods. This is not particularly surprising, as we would expect that, for example, if a sufficiently important positive or negative news about a stock is released that large price movements would quickly ensue in an efficient market.

⁹In doing so, overlapping *portfolios* are formed rather than returns. Thus, assuming no autocorrelation in the returns on the momentum portfolios, it is not necessary to correct for serial correlation.

4.4 The Returns of Intraday Winner and Loser Portfolios

This section presents the results of a trading strategy based upon forming portfolios conditional on performance over a short time horizon of 10-60 minutes. By going long winners and short losers, it is possible to see whether a reversal or relative strength effects dominates intraday.

Since the data are filtered to produce 5-minute intraday returns, it is convenient to use 5-minute returns in the analysis. Table 4.2 reports the average 5-minute returns for the relative strength portfolios. For each formation period ($J=10, 20, 30$ and 60 minutes) and holding period ($K=10, 20, 30$ and 60 minutes), the average returns on the buy portfolio, sell portfolio and the (zero cost) buy-sell portfolio are shown. For example, the third row of the fourth column gives the mean 5-minute percentage return (0.000056, i.e. 0.0056%) for selling the portfolio of losers over the previous 10 minutes and holding this portfolio for 20 minutes.

Looking down the table shows the four portfolio formation periods (10, 20, 30 and 60 minutes), whilst the columns show the four holding periods for these portfolios (10, 20, 30 and 60 minutes). The third column shows that all of the portfolio returns for a holding period of 10 minutes ($K=10$) are statistically significant at the 5% level. The buy portfolio represents the return from buying the decile of stocks that have performed best over the formation period of J minutes. The results show that these returns are all *negative*; a strategy of buying 'winners' is therefore not profitable. In fact, this shows that the success of winners in the formation period reverses and thus contrarian effects are present, i.e. it is profitable to *sell* winners. This provides initial confirmation of a reversal effect intraday.

The sell portfolios show the returns from selling the decile of stocks that were the poorest performers in the previous J minutes. All of these returns are positive

Table 4.2: Returns of relative strength portfolios

The sample period is January 1 2005 to December 31 2005 comprising the constituents of the S&P 500 index (adjusted monthly for additions and deletions to the index). The relative strength portfolios are formed in the manner of Jegadeesh and Titman (1993), i.e. they are formed on J -minute lagged returns and held for K -minutes. The values for J are listed in the first column and K in the first row. S&P constituent stocks are ranked based on their J -minute lagged returns, in ascending order. The *sell* portfolio is the equally weighted portfolio of stocks in the lowest past return decile. The *buy* portfolio is the equally weighted portfolio of stocks in the highest returns decile. The mean 5-minute return for the portfolios are presented in this table. p -values are presented in parentheses.

$J =$	$K =$	10	20	30	60
10	Buy	-0.000166 (0.0004)	-0.000105 (0.0004)	-0.000081 (0.0016)	-0.000067 (0.0147)
10	Sell	0.000126 (0.0021)	0.000056 (0.1818)	0.000036 (0.0587)	0.000021 (0.3740)
10	Buy-sell	-0.000291 (0.0004)	-0.000161 (0.0010)	-0.000117 (0.0017)	-0.000088 (0.0030)
20	Buy	-0.000143 (0.0013)	-0.000081 (0.0035)	-0.000066 (0.0059)	-0.000053 (0.0211)
20	Sell	0.000118 (0.0120)	0.000057 (0.1164)	0.000044 (0.1272)	0.000020 (0.4920)
20	Buy-sell	-0.000261 (0.0025)	-0.000138 (0.0172)	-0.00010 (0.0151)	-0.000073 (0.0533)
30	Buy	-0.000137 (0.0044)	-0.000081 (0.0093)	-0.000072 (0.0219)	-0.000050 (0.0479)
30	Sell	0.000118 (0.0095)	0.000070 (0.0315)	0.000052 (0.0716)	0.000021 (0.4846)
30	Buy-sell	-0.000256 (0.0044)	-0.000151 (0.0177)	-0.000125 (0.0197)	-0.000071 (0.1084)
60	Buy	-0.000134 (0.0125)	-0.000080 (0.0372)	-0.000068 (0.0665)	-0.000043 (0.1684)
60	Sell	0.000130 (0.0129)	0.000073 (0.0838)	0.000054 (0.1734)	0.000018 (0.6757)
60	Buy-sell	-0.000265 (0.0097)	-0.000154 (0.0443)	-0.000121 (0.0823)	-0.000061 (0.3567)

(and significant at the 1% level). It is therefore profitable to buy the 'loser' portfolio, rather than sell as suggested by the momentum trading strategy, again lending support to the dominance of reversal at short time horizons.

The returns to the zero cost buy-sell portfolio for a holding period, K , of 10 minutes are significant for all formation periods at the 1% level. Whilst momentum would not form a profitable trading strategy, these results show that there is a reversal effect present. Selling 'winners' and buying 'losers' over the previous 10-60 minutes and holding them for 10 minutes is a profitable strategy.

Having established that over the shortest time horizon of 10 minutes returns tend to reverse, we extend the holding period to 20 minutes ($K=20$). Again, these results show that a contrarian/reversal effect rather than momentum effect is present. The mean 5-minute returns from buying the portfolio of the best performers are all negative, with these returns all significant at the 5% level. All of the sell portfolio returns are positive, although only the 30 minute formation period ($J=30$) is statistically significant. The zero-cost buy-sell portfolio produces a statistically significant return for all formation periods.

All of the buy returns for a 30 minute holding period ($K=30$) are again negative and significant at the 5% level, with the exception of portfolios formed over the longest time of 60 minutes. Again, all sell portfolio returns are positive albeit not statistically distinguishable from zero. Buy-sell portfolios formed on the basis of returns over the preceding 10, 20 and 30 minutes produce a significant negative return.

The longest holding period ($K=60$) results in the last column show that forming zero-cost winners-losers portfolio over 20, 30 and 60 minutes and holding it for 20, 30 or 60 minutes does not produce a significant return. However, these results show that prices change rapidly over a shorter time horizon but revert over a longer period. This is seen as the buy-sell portfolio return for a portfolio formed on the basis of returns over the previous 10 minutes and held for an hour ($J=10/K=60$)

shows a significant negative return.

These results clearly show that a contrarian strategy is profitable at very short (intraday) time horizons; significant reversal, and not momentum, effects are present. With a formation and holding period of up to 30 minutes a highly significant reversal effect is present, with the implication that buying 'losers' and selling 'winners' is profitable. Figure 4.1 presents a graphical representation of the mean 5-minute returns across the possible combinations of formation and holding periods. This gives a better idea of the size of the returns that are generated. It is seen that the size of the mean 5-minute returns over the 10 minute holding period ($K=10$) are much larger than for the longer holding periods. This suggests that the reversal of intraday prices occurs relatively quickly. Likewise, the mean 5-minute return is lower for $K=30$ as opposed to $K=20$ and $K=60$ versus $K=30$. This is an important result, and suggests that reversal effects are comparatively short lived. If some market participants have a tendency to overreact, prices return to a 'fair' level relatively quickly. This is consistent with the large number of day traders and professional traders in the market seeking to exploit short-term profit opportunities.

The results show that for all portfolio formation periods, as the holding period (K) increases, profits are monotonically decreasing. This demonstrates that reversal effects decay relatively quickly. In contrast, holding K constant and looking at the different formation periods does not produce a discernible pattern between $J=10$, 20, 30 and 60 minutes. These results do, however, show that one can condition portfolios based on relative returns up to one hour previously and hold for 10 and 20 minutes, and still observe a subsequent reversal effect. The implications for traders are that stock prices do exhibit an intraday reversal effect, with profits shown over times of up to an hour in the trading day. It is reasonable to suppose that this strategy may be exploitable by day traders who enjoy low transactions costs.




Figure 4.1: Returns of relative strength portfolios

It is now commonly recognised that many anomalies seem to be sensitive to particular time periods. For example, the January effect, of which possible explanations include the publication of accounting news (Rozeff and Kinney, 1976), and tax-loss selling Reinganum (1983). There is also evidence that investors' behaviour is conditional upon the month of the year (Ritter, 1988). Given the wide range of seasonal anomalies documented in the literature, it is sensible to investigate whether the reversal effect documented above varies by calendar month. The next section decomposes the results by month to assess the robustness of intraday reversal profitability.

4.4.1 Subperiod analysis

Table 4.3 reports the average returns of the zero-cost buy-sell portfolios by calendar month. Four symmetrical trading strategies are presented: $J=10/K=10$, $J=20/K=20$, $J=30/K=30$ and $J=60/K=60$. These are seen in the column headings looking across the table. Looking down the table shows the calendar months of the year. Figure 4.2 presents a graphic representation of these results.

The returns reported from the shortest symmetrical formation/holding period ($J=10/K=10$) show that reversal, rather than momentum, effects consistently prevail. Buy-sell portfolio returns are significant, and positive, for all months, with the exception of July and October. All of the buy-sell returns are negative—employing a reversal based trading strategy is profitable. Looking at January in particular, there is no apparent difference in the sign or magnitude of the portfolio return compared with other months in the year.

Increasing the formation and holding periods leads to fewer of the buy-sell portfolios providing significant returns. At $J=20/K=20$ all of the average returns are again negative with the exception of September. In this case, a positive return is seen on the buy-sell portfolio, suggesting a conventional momentum strategy

Table 4.3: Returns of relative strength portfolios by calendar month

The sample period is January 1 2005 to December 31 2005 comprising the constituents of the S&P 500 index (adjusted monthly for additions and deletions to the index). The relative strength portfolios are formed on J -minute lagged returns and held for K -minutes. The mean 5-minute return for the portfolios are shown. For example, column three shows returns for portfolios formed on 20 minute lagged returns and held for 20 minutes. The table reports the average return of the zero-cost buy-sell portfolio by month; January through to December are shown individually in the rows of the table. p -values are presented in parentheses.

	J=10/K=10	J=20/K=20	J=30/K=30	J=60/K=60
Jan.	-0.000250 (0.0022)	-0.000133 (0.0051)	-0.000117 (0.0497)	-0.000028 (0.6732)
Feb.	-0.000236 (0.0022)	-0.000164 (0.0063)	-0.000097 (0.0819)	-0.000032 (0.5861)
Mar.	-0.000295 (0.0010)	-0.000159 (0.0070)	-0.000108 (0.0558)	-0.000068 (0.2793)
Apr.	-0.000302 (0.0040)	-0.000225 (0.0017)	-0.000149 (0.0199)	-0.000040 (0.5259)
May	-0.000334 (<.0001)	-0.000197 (<.0001)	-0.000113 (0.0383)	-0.000091 (0.1265)
June	-0.000227 (0.0027)	-0.000111 (0.0180)	0.000069 (0.2388)	-0.000086 (0.0929)
July	-0.000177 (0.1625)	-0.000105 (0.1381)	-0.000033 (0.7284)	-0.000069 (0.4237)
Aug.	-0.000264 (0.0025)	-0.000159 (0.0016)	-0.000105 (0.0536)	-0.000028 (0.6915)
Sept.	-0.000219 (0.0184)	0.000125 (0.0319)	-0.000081 (0.1390)	-0.000080 (0.2497)
Oct.	-0.000612 (0.0563)	-0.000180 (0.0128)	-0.000215 (0.0887)	0.000133 (0.6083)
Nov.	-0.000305 (0.0007)	-0.000231 (<.0001)	-0.000163 (0.0459)	-0.000094 (0.3190)
Dec.	-0.000368 (0.0072)	-0.000180 (0.0042)	-0.000110 (0.0681)	-0.000121 (0.1131)




Figure 4.2: Buy-sell portfolio returns by calendar month

Table 4.4: Returns of relative strength portfolios by day of week

The sample period is January 1 2005 to December 31 2005 comprising the constituents of the S&P 500 index (adjusted monthly for additions and deletions to the index). The relative strength portfolios are formed on J -minute lagged returns and held for K -minutes. The mean 5-minute return for the portfolios are shown. For example, column three shows returns for portfolios formed on 20 minute lagged returns and held for 20 minutes. The table reports the average return of the zero-cost buy-sell portfolio by day of the week, with Monday through to Friday shown by row. p -values are presented in parentheses.

	J=10/K=10	J=20/K=20	J=30/K=30	J=60/K=60
Monday	-0.000285 ($<.0001$)	-0.000147 (0.0045)	-0.000097 (0.0607)	-0.000092 (0.1763)
Tuesday	-0.000300 (0.0013)	-0.000280 (0.8111)	-0.000131 (0.0663)	-0.000023 (0.7781)
Wednesday	-0.000284 (0.0018)	-0.000174 (0.0168)	-0.000113 (0.0779)	-0.000003 (0.9732)
Thursday	-0.000282 (0.0008)	-0.000200 (0.0010)	-0.000163 (0.0076)	-0.000117 (0.1136)
Friday	-0.000258 (0.0024)	-0.000160 (0.0013)	-0.000085 (0.0760)	-0.000059 (0.2831)

can be profitably employed. This is an anomaly in the results; all other months for $J=20/K=20$ show a significant negative return.

For $J=30/K=30$, the negative return on the buy-sell portfolio is significant only in January, April, May and November (a positive but insignificant return is seen in June). For the longest formation and holding period $J=60/K=60$ all buy-sell portfolios show a negative return with the exception of October. However, these returns are insignificant at the 5% level (July is significant at 10%).

Taken together, these results show that dis-aggregating the study by month confirms the dominance of a reversal/contrarian effect. The profits from such a strategy appear to be consistent throughout the calendar year.

4.4.2 Day of week and intraday relative strength portfolios

Existing evidence shows that there are systematic patterns in stock returns over the trading week (for example French, 1980). Given our knowledge of this, it is interesting to establish if the reversal effect is conditional upon the day of the week. Table 4.4 shows the mean 5-minute returns of the zero-cost buy-sell portfolios by weekday. As above, results for four symmetrical trading strategies of 10, 20, 30 and 60 minutes are presented across columns. The rows show the returns by day of the week. It is clear that for the shortest formation/holding period of $J=10/K=10$, that the mean returns are all negative and significant. Figure 4.3 gives a visual indication of the returns broken down by day. It can be seen that mean returns are of broadly similar magnitude throughout the week, although marginally lower at the end than the start.

The $J=20/K=20$ strategy displays significant mean returns on all weekdays apart from Tuesday. All returns are negative. For $J=30/K=30$ all mean returns are again negative, although only the return for Thursday is significant at the 5% level. The mean returns for $J=60/K=60$ are generally small in size. None of these returns proved to be statistically significant.

These results show that the intraday momentum strategy does not produce markedly different results based on the day of the trading week. Lakonishok and Maberly (1990) document greater propensity for individual investors to sell on a Monday, whilst institutions trade very little, leading to low volume. This selling pressure, however, does not influence the intraday reversal effect.

4.4.3 Time of day and intraday relative strength portfolios

Considerable evidence has been presented in the literature to suggest that there is an intraday U-shaped curve in stock prices, volume and volatility (Wood et al., 1985; Harris, 1986; Foster and Viswanathan, 1990). This means that more market activity takes place at the start and end of the trading day, near to the open and close of the market. In examining the returns of an intraday relative strength based




Figure 4.3: Buy-sell portfolio returns by day

Table 4.5: Returns of Relative Strength Portfolios by Hour

The sample period is January 1 2005 to December 31 2005 comprising the constituents of the S&P 500 index (adjusted monthly for additions and deletions to the index). The relative strength portfolios are formed on J -minute lagged returns and held for K -minutes. The mean 5-minute return for the portfolios are shown. For example, column three shows returns for portfolios formed on 20 minute lagged returns and held for 20 minutes. The table reports the average return of the zero-cost buy-sell portfolio by hour, with seven hourly subperiods of the trading day shown individually. p -values are presented in parentheses.

	J=10/K=10	J=20/K=20	J=30/K=30	J=60/K=60
9:30-10:30	-0.000196 (0.0200)	-0.000183 (0.0252)	-0.000183 (0.0247)	-0.000183 (0.0247)
10:30-11:30	-0.000073 (0.0001)	-0.000074 (0.0043)	-0.000086 (0.0008)	-0.000076 (0.0020)
11:30-12:30	-0.000027 (0.2782)	0.000195 (0.3164)	0.000038 (0.0606)	-0.000034 (0.1127)
12:30-13:30	-0.000103 (0.0010)	-0.000055 (0.0241)	-0.000039 (0.1083)	-0.000047 (0.0370)
13:30-14:30	-0.000080 (0.0301)	-0.000028 (0.3743)	0.000006 (0.8258)	0.000044 (0.0321)
14:30-15:30	-0.000116 (0.0002)	-0.000024 (0.3950)	0.000008 (0.7900)	0.000050 (0.0875)
15:30-16:00	-0.001342 (0.0045)	-0.000665 (0.0388)	-0.000627 (0.0245)	-0.000331 (0.3274)

trading strategy, it is therefore desirable to dis-aggregate results and investigate how these vary within the trading day. Should profits be driven by greater returns at specific times of the day, there are clear implications for traders looking to profit from intraday reversal effects.

Table 4.5 reports average returns for the zero cost buy-sell portfolio for six discrete hourly periods from 9:30 to 15:30 (and a half hour from 15:30 to 16:00). Looking across the table, four symmetrical trading strategies are presented: $J=10/K=10$, $J=20/K=20$, $J=30/K=30$ and $J=60/K=60$, and looking down the table returns are separated into the hourly subperiods. Given the U-shaped pattern of activity over

the course of the trading day, it makes sense to evaluate the subperiods of one hour after the opening and one hour before the closing first, when activity is likely to be greatest. For the opening hour 9:30-10:30, the mean return of the buy-sell portfolio is negative for all four of the J/K combinations. These returns are significant, and suggest a broad reversal effect is present after the market opens.

The $J=10/K=10$ results column shows negative buy-sell portfolio returns throughout the trading day. All of these results are significant with the exception of the 11:30-12:30 subperiod. This result may be due to the lower volume that is seen in the middle of the day. This accords with the notion of a U-shaped pattern to market activity, with lower volumes in the middle of the day. In this case, it is possible that fewer trades are undertaken by intraday traders, thus there is less of a tendency for overreaction leading to a reversal effect.

Moving to $J=20/K=20$, the buy-sell returns are again negative for all periods apart from 11:30-12:30, which exhibits an insignificant positive return. The negative returns for the subperiods apart from 11:30-12:30 are statistically significant at the 5% level for 9:30-10:30, 10:30-11:30, 12:30-1:30 and 15:30-16:00. Extending the holding and formation periods to 30 minutes, the results for $J=30/K=30$ sees the first two hours and the last half hour of the trading day producing significant negative buy-sell returns. However, between 11:30 and 15:30 the returns are insignificant. Apart from 11:30-12:30 all of the hours see a positive but insignificant return. At the longest $J=60/K=60$ symmetrical formation/holding period, the first two hours of the data from significant negative buy-sell returns on the zero-cost portfolio. However, for 13:30-15:30 two hourly portfolio mean returns are shown to be significantly positive. The last half hour of the trading day is the only one presented not to show a significant negative return.

These results show that the reversal effect seen in the previous sections is not present uniformly across the trading day. A significant buy-sell return is seen over the first hour of the trading day for all J/K combinations. Apart from $J=60/K=60$




Figure 4.4: Returns of Relative Strength Portfolios by hour

we see a significant negative buy-sell return in the last half hour of the trading day. During the trading day, the picture is more mixed, and both insignificant portfolio returns and significant positive portfolio returns are exhibited. Figure 4.4 presents these results graphically and helps to convey the relative magnitude of these returns. It is immediately clear that the magnitude of the buy-sell portfolio returns is greater at the start and end of the trading day. For instance, panel b shows the returns for $J=20/K=20$. For the first hour and last half hour of trading, the mean 5-minute return on the buy-sell portfolio is -0.0183% and -0.0665% respectively. By contrast, the (statistically significant) return for 12:30-13:30 is -0.000028%. A similar pattern is seen in the other three symmetrical formation/holding periods. This result reinforces the conclusion that the reversal effect is not constant over the trading day. It is consistent with non-uniform market activity over trading hours. One possible explanation exists in the already well known U-shaped pattern in volume and volatility. The implication for traders is that the start and end of the day are the most profitable periods for an intraday reversal strategy.

4.4.4 Size and intraday relative strength portfolios

Whilst the sample of S&P 500 stocks broadly represents the largest and most actively traded stocks, there is still considerable variation between the constituents in terms of size and trading volume. It is therefore interesting to investigate if market capitalisation affects momentum profitability. Table 4.6 presents the results of dividing the sample into equal quintiles based on market capitalisation. Size quintile 1 represents the largest companies and quintile 5 the smallest.

Columns two to four of Table 4.6 show the previously used four symmetrical formation and holding periods. Column two shows the mean 5-minute returns for the $J=10/K=10$ buy-sell portfolio over the five quintiles. All of these returns are significant at the 1% level. The buy-sell portfolio is therefore significantly profitable

Table 4.6: Returns of relative strength portfolios by market capitalisation

The sample period is January 1 2005 to December 31 2005 comprising the constituents of the S&P 500 index (adjusted monthly for additions and deletions to the index). The relative strength portfolios are formed on J -minute lagged returns and held for K -minutes. The mean 5-minute return for the portfolios are shown. For example, column three shows returns for portfolios formed on 20 minute lagged returns and held for 20 minutes. The table reports the average return of the zero-cost buy-sell portfolio with quintiles determined by market capitalisation. p -values are presented in parentheses.

Size Quintile	J=10/K=10	J=20/K=20	J=30/K=30	J=60/K=60
1	-0.000339 (0.0001)	-0.000195 (0.0006)	-0.000132 (0.0170)	-0.000079 (0.2597)
2	-0.000246 (0.0018)	-0.000160 (0.0061)	-0.000128 (0.0246)	-0.000088 (0.2008)
3	-0.000238 (0.0034)	-0.000096 (0.2085)	-0.000059 (0.4282)	-0.000053 (0.4818)
4	-0.000297 (0.0039)	-0.000114 (0.0633)	-0.000102 (0.0601)	-0.000071 (0.2795)
5	-0.000298 (0.0006)	-0.000106 (0.1241)	-0.000130 (0.0059)	-0.000012 (0.8592)

for a reversal strategy, irrespective of the size group. Moving to $J=20/K=20$ all of the buy-sell returns are again negative; however, only those for the two largest quintiles of stocks prove to be statistically significant. For the longer $J=30/K=30$ formation/holding period these two largest quintiles again show significant negative returns yet the negative return for the fifth quintile is also significant. None of the negative returns reported for the $J=60/K=60$ formation/holding period prove significant.

Figure 4.5 more clearly represents the magnitude of these returns. It is shown that—in common with previous results—that the mean 5-minute returns decrease with the extension of the formation/holding period. It is interesting to note that the highest average mean returns for the buy-sell portfolios are for the quintiles 1 and 2, representing the largest stocks. At the shortest formation/holding period of $J=10/K=10$ the mean returns are markedly higher for the largest size quintile. These results suggest that the more actively traded larger stocks present a better opportunity for counter-momentum profitability.

4.5 Conclusions

This chapter tests for the presence of intraday momentum and reversal effects, and forms portfolios based on prior relative returns of 10-60 minutes, using high-frequency trade data for S&P constituents throughout 2005. In forming relative strength portfolios, in a similar manner to Jegadeesh and Titman (1993), and buying winners and selling losers, it is found that a reversal effect—rather than momentum—is seen intraday. The size of the reversal effect is greatest 10-minutes after portfolio formation and decays thereafter. This is an important result for both academics and market professionals. This latter group includes a substantial number of day traders, which survey evidence shows are largely momentum traders (Mizrach and Weerts, 2009). This being the case, these results show that




Figure 4.5: Buy-sell portfolio returns by market capitalisation

such traders may be misguided, and instead should look towards identifying reversal trades.

The momentum effect is well established at time horizons measured in months, but there has previously been almost no detailed investigation of whether momentum is seen within the trading day. In addition, the evidence of intraday reversal and contrarian effects is inadequate. Several papers have sought to look at intraday reversals, but are generally confined to an event study methodology looking at the tendency of price to reverse after large changes at the open and/or using a limited data set. This study uses the established method of constructing portfolios based on momentum over various formation periods and holding these portfolios for a range of different holding periods, between 10 and 60 minutes. In doing so, the presence of an intraday reversal effect has been identified.

The returns of the zero-cost buy-sell portfolio, comprised of winners minus losers, were also obtained by day of the week, month, time of day and by firm size. It was found that it was not possible to isolate the reversal effect to a particular month or day. However, and in accordance with existing evidence showing a U-shaped pattern in volume and volatility throughout the trading day, with increased activity around the open and close, mean returns were shown to be largest at the start and end of the day. Separating the constituents of the S&P500 into size quintiles showed that the reversal effect was most pronounced amongst the two largest size quintiles. These stocks possess relatively greater trading volumes.

The analysis in this chapter has shown the presence of an intraday reversal effect. Market participants seeking to exploit this would need to place frequent orders which would remain open for only a short amount of time in order to achieve a relatively small profit. This naturally raises the question of transaction costs which, as discussed previously, have been shown to void the profitability of many technical trading strategies. In examining the profitability of traders who operate at the shortest time horizons, Mizrach and Weerts (2009) examine a large

number of day traders and find that 55% make a profit net of transaction costs. This may seem unlikely given one-way costs such as the figures of 0.18% (Jones, 2002) and 0.23% (Berkowitz et al., 1988). A possible mediator here is that many day trading firms offer flat-fee commissions. For instance, Jordan and Diltz (2003) study trading activity at a firm with a base commission of \$14 per trade, with the addition of a fee from the ECN used.¹⁰

In relevance to an intraday technical trading strategy, both Mizrach and Weerts (2009) and Jordan and Diltz (2003) show that a significant portion of day traders are able to make a profit net of transaction costs. This suggests that such traders may be able to take advantage of a short time horizon reversal strategy. However, it is highly desirable that further research investigates this issue by looking to account for transaction costs on a trade-by-trade basis. This would afford an insight into the profitability of an intraday reversal strategy for market participants including both day traders and institutions.¹¹

It is also important to consider the potential impact of short sale costs and constraints. Whilst future research could consider these more explicitly, this study focusses on large S&P 500 constituents. D'Avolio (2002) notes that "general collateral" stocks, which are the easiest stocks to borrow for shorting purposes, have a mean value-weighted cost of 0.17% per annum. It is noted that, as S&P 500 stocks are held in large quantities by passive investment vehicles (specifically index trackers), that these "are almost always general collateral" (D'Avolio, 2002, p.273), and thus costs are much lower.

Given the very short time horizons investigated in this study, it seems unlikely that profitability can be related to time varying risk premia. This argument is presented by Marshall et al. (2008b) related to an examination of five intraday

¹⁰These costs are between a flat \$0.50 and \$2.50 for the main ECNS and \$0.015 per share for other ECNs.

¹¹Some initial empirical work beyond this thesis has been undertaken. This research looks at forming winner and loser portfolios on smaller portfolios by using deciles, with NYSE TAQ data for 2008. Early results concur with those presented above.

technical trading rules. However, one possibility for future research is to look at the drawdown from an intraday relative strength strategy.

The evidence presented in this study shows that there is a pronounced intraday reversal effect. In reconciling the relative strength strategies seen to be employed by fund managers and existing academic evidence that a reversal effect was prevalent, Jegadeesh and Titman (1993) showed the importance of the time horizon over which portfolios are formed and evaluated. The results show that shortening the time horizon as far as possible with high-frequency intraday data gives rise to a contrarian rather than momentum effect. This result is consistent with the notion that market participants overreact to information even at the shortest time-horizons.

Chapter 5

Point and Figure Trading

5.1 Introduction

Convincing evidence shows that technical analysis is actively employed by market participants to make investment decisions (Taylor and Allen, 1992; Oberlechner, 2001; Gehrig and Menkhoff, 2006). A large number of studies have demonstrated that technical trading rules can generate profits; for example, Brock et al. (1992) for the widely-used moving average. Earlier chapters in this thesis have shown the economic value of the head and shoulders pattern and the existence of intraday reversal. However, while an array of well-known and easily tested rules have been evaluated, there are still important areas of technical analysis that have not been examined. This chapter investigates perhaps the oldest form of technical analysis: point and figure charting. As yet, an extremely limited amount of work has been undertaken into the profitability of a technical trading strategy based on point and figure.

As a type of technical analysis, point and figure charting is unique in plotting price data independent of time. The method has well defined rules for constructing charts, designed to isolate important price moves and use these to establish trading signals. The technique of point and figure has been used by traders for over 100 years. Murphy (1999) notes that point and figure was previously known as the “book method”, a term coined by Charles Dow in a *Wall Street Journal* editorial in 1901. Murphy determines that Dow indicated the method had been in use for around 15 years at the time, meaning it dates from the mid-1880s. The earliest detailed account of point and figure is provided by DeVilliers and Taylor (1933), who present a thorough account of the methodology, giving clear rules, applications and examples for practitioners. Crucially, however, this technique is still in active use today, and there are plentiful examples of literature on point and figure analysis aimed at traders (for example Du Plessis, 2005; Dorsey, 2007; Weber and Zieg, 2003; Dorsey et al., 2007). Furthermore, point and figure charts can be produced by almost all professional trading software and many stock charting

web sites. This long history is important given that one of the greatest sources of academic scepticism concerning technical analysis is data mining. A set of clear trading rules—seldom present in other types of technical analysis—and the long-standing knowledge and use of point and figure charting means that this is a particularly valuable and interesting area of study.

In spite of this, the existing academic literature provides only a very limited investigation of the profitability of trading strategies based upon point and figure charting, and is subject to a number of important limitations. First, we know very little about the performance of a point and figure trading strategy for a substantive sample of stocks. Studies have tended to concentrate on a single futures contract, single stock index or, in a limited number of cases, a small sample of stocks.¹ Second, virtually all current research focusses on using daily or weekly closing prices. This is important as point and figure was originally intended for use with ‘real-time’ data by floor traders. Whilst the technique has come to also be employed at daily (or longer) time horizons, the practitioner texts still clearly advocate its use on ultra high-frequency data (for example, Dorsey, 2007). A recent study by Anderson and Faff (2008) attempts to address this limitation to some extent. However, in solely investigating the S&P 500 futures contract from 1990 to 1998, we still do not have any evidence on the profitability of intraday point and figure trading rules for individual stocks.

The subject of point and figure charting is also an interesting area to examine in relation to earlier work in this thesis. The first two empirical chapters looked at the head and shoulders formation. The methodology called first for the identification of localised maxima and minima prior to these points being used to identify head and shoulders patterns. However, there is some degree of ambiguity concerned the specification of head and shoulders patterns in the practitioner literature. By contrast, point and figure trading signals—which, as will be seen below, are

¹However, as will be shown below, existing studies looking at stocks suffer from important limitations and none use intraday data.

determined by pattern-like formations on point and figure charts—are very well defined, with clear specifications.

In this study, intraday data from the New York Stock Exchange Trade and Quote database is used. The sample spans all of 2005, and includes all S&P 500 constituents, re-sampled monthly to take account of additions and deletions from the index. The approach taken is to evaluate point and figure trading rules for a large sample of stocks with this ultra high-frequency data. As such, this chapter makes an important contribution to our knowledge of technical analysis in four main ways. First, evaluating the profitability of such long-standing and established trading rules is valuable as, unlike virtually all other forms of technical analysis, data mining is not a central consideration. Second, point and figure charting is an important area of technical analysis that, as yet, has been significantly under-investigated in the literature. Third, the small amount of existing research is limited in a number of respects. Evaluating point and figure using ultra high-frequency data for a large sample of stocks serves to address these limitations and advance our knowledge. Fourth, as point and figure is a trading strategy widely employed by market professionals, the results have considerable practical relevance.

The remainder of this chapter is organised as follows. Section 5.2 presents a survey of the literature on point and figure trading strategies. This serves to further highlight the need for a comprehensive study, which forms the basis of this chapter. In addition, a detailed exposition of the nature of point and figure charts and trading signals is given. Section 5.3 describes the ultra high-frequency data used, and presents the methodology employed. Section 5.4 contains results and a discussion of the point and figure trading strategies. Finally, section 5.5 provides concluding remarks and suggestions for subsequent research.

5.2 Point and Figure Analysis and Literature

5.2.1 Construction of point and figure charts

It is important to be clear on what differentiates the technique of point and figure from other technical analysis methods, as these factors also serve to make this a particularly interesting topic for study. Point and figure charts are constructed on a grid, by plotting a series of 'X's to represent upward price movements and 'O's for downward price movements. These 'X's and 'O's are stacked in columns, which visually represent the magnitude of price moves over time. However, the charts themselves differ from all other forms of technical analysis chart by not reflecting time in a linear fashion. Demonstrating this by way of an example, Figure 5.1 shows a randomly generated series of 90 prices for a stock, plotted as a conventional line chart. Figure 5.2 shows the same prices on a point and figure chart.




Figure 5.1: Standard line chart of stock price

The point and figure chart has 25 data points, compared to the 90 observations in the original series. The point and figure technique has effectively filtered the 'noisy' data, clearly isolating the main price moves over the period. This is achieved through relatively simple means. Individual points plotted as 'X's or 'O's on the




Figure 5.2: Point and figure chart of the same series as Figure 5.1, with a box size of \$1 and reversal of 3 boxes

chart only appear when the price rises or falls by a set amount; this is referred to as the 'box size'. Thus, small price changes are disregarded. This element of point and figure charting is therefore similar to a filter rule. However, unlike a filter rule, which does not alter the display of prices in a chart, a key distinguishing feature of point and figure charting is how these filtered price moves are arranged into columns. Figure 5.2 displays six columns, alternating between 'X's and 'O's. A detailed methodology for construction of point and figure charts is given below; however, it is worth recognising now that when price rises (falls) more than a set amount, an 'X' ('O') is plotted. Dorsey (2007) asserts that columns of 'X's and 'O's represent times when demand overwhelms supply and supply outstrips demand, respectively.

A visual comparison of the original series to the point and figure chart shows that the columns of 'X's and 'O's have succeeded in isolating the main price swings. The individual columns thus exhibit the cumulative price change for each of these swings. As such, the way in which an existing column is ended, followed by a rightwards movement and the initiation of a new column, is of crucial importance. For this to take place, the point and figure charting technique requires the price

to move against the prevailing trend by a set multiple of the box size (price filter), referred to as the 'reversal amount'. Again, this serves to disregard unimportant whipsaws in price and isolate important moves. This filtering mechanism is very useful for building trading strategies. Such strategies can then be seen as only responding to important price movements and can use the trends that are identified by point and figure charts. By constructing charts in this way, traders can also employ well-defined rules to initiate buy and sell trades. This is in contrast to technical analysis methods such as the moving average, where a wide variety of average lengths could be chosen, potentially resulting in data mining. Other types of technical analysis do not possess this inbuilt advantage, which makes point and figure unique.

Two key features of point and figure are the box size and reversal amount. The box size represents the unit of measurement on the chart. Individual price movements at a point in time or cumulative price moves that are smaller than the box size are ignored. If we denote the box size as B , then for a new box to be filled on the chart the condition is $P_t - P_{t-n} > B$. The 'reversal amount' constitutes the number of boxes that have to be accumulated—against the direction of the prevailing column—for the current column to be 'reversed' and a new column initiated. Denoting the reversal amount as R , the current column as C_j , and assuming $R = 3$, if the current column is 'X's, we require $P_t < (C_j^{high} - 3B)$ to move rightwards on the chart and start a new column of 'O's. For example, the standard 3-box chart implies that if the box size is \$1 then a counter-move of greater than or equal to \$3 is needed for a new column. Together, these two features of point and figure charts lead to one of the main advantages of the technique: the chart forms a filtered representation of price moves. Thus, point and figure charts serve to provide a clearer representation of the important price moves and trends in 'noisy' price data.

It is easiest to appreciate the nature of point and figure charts by example.

Suppose that we are working with a box size of \$1 and a standard 3-box reversal chart. Assume a security trades at \$50.40 at time t . First, one identifies whether the current 'column' is composed of 'X's or 'O's. Assume that an uptrend is in place and, therefore, that the current column is of 'X's. If the existing high of this column is \$50, as plotted on the chart, the technical analyst would ignore the price of \$50.40 at t (the 40¢ difference is smaller than the box size of \$1) and look to the next time period. At $t + 1$ the security trades at \$54. If the box size is \$1 then this represents a rise of four boxes from the previous high of the column of 'X's on the chart. The analyst would then mark on three further 'X's to take the high of the column up to \$54.

At $t + 2$ the price falls by \$2 to \$52. Recall that this is a 3-box chart with a \$1 box size. This implies that the price needs to change by the equivalent of $3 \times \$1$ to start a new column: the price would need to fall by \$3 for the current column of 'X's to be abandoned. Since the price at $t + 2$ has fallen by only two boxes, or \$2, then the technical analyst does not make any changes to the chart. No new plots are made and the column high remains at \$54.

At $t + 3$ the price change is +\$1 to \$54. Again, the technical analyst makes no change to the chart as \$54 is still equal to the highest box of 'Xs' on the chart (\$54 at $t + 1$). However, at $t + 4$ the security's price falls by \$4 to \$50. As this fall is greater than 3 boxes, this price change signals the end of the current column of 'X's and the start of a new column of 'O's. Accordingly, the technical analyst advances right one column and enters an 'O' one box below the highest box containing an 'X', in this case \$53. 'O's are then filled downwards to the current price of \$50, i.e. three 'O's in total.

A change of -\$1 to \$49 occurs at At $t + 5$. The analyst logs this on the chart with a single 'O' appended to the bottom of the current column of 'O's. The reversal amount is symmetric. So, for example, if the price of the security rallies to \$53 at $t + 6$ then the analyst would shift right one column to begin a new column of

'X's. Starting at \$50 (one box above the previous column low) three 'X's are drawn in the new column.

In this fashion, a point and figure chart is built. By ignoring periods where price is unchanged, a price change in the opposite direction to the column smaller than the reversal amount, and appending rises (falls) to an existing column of 'X's ('O's), a visual record of the stock price movement is built which does not show time linearly. In contrast to a conventional price chart, all of the 'X's and 'O's can be viewed as observations of interest to the analyst.

To illustrate how the point and figure method works with a real example we can study the intraday price chart for Microsoft Corporation, an S&P 500 stock traded on the New York Stock Exchange, for 3 January, 2005. Figure 5.3 shows the raw price series for the day. It can be seen that the chart is very 'noisy', and it is somewhat difficult to establish clear trends. By contrast, Figure 5.4 shows a point and figure chart for the same data. This is a standard 3-box reversal chart, with a box size of \$0.02, leading to 33 columns being plotted. By using point and figure, it is immediately clear visually that the chart depicts the important price moves that took place during the trading day. As well as allowing trading rules to be employed (see below), this also affords traders an important insight into levels of support and resistance. These are central concepts to all aspects of technical analysis, yet their identification is often subjective. This is not the case with the point and figure technique.

However, this example can also serve to demonstrate why the box size and reversal amount are so important in point and figure charting. Figure 5.5 uses the same data but with a 5-box (rather than 3-box) reversal amount. The box size remains at \$0.02. The effect is that a larger magnitude of price reversal against the prevailing trend is required before one column is ended, and a new one initiated. The net result is that the point and figure chart representing the same set of intraday data now has only 14 columns (rather than 33). Arguably, this provides an even




Figure 5.3: Intraday price chart of Microsoft Corporation - 3 January, 2005




Figure 5.4: Point and figure chart of stock price shown in Figure 5.3 with a box size of \$0.02 and reversal of 3 boxes

clearer picture of the intraday price movements and isolates the major trends.




Figure 5.5: Point and figure chart of stock price shown in Figure 5.3 with a box size of \$0.02 and reversal of 5 boxes

Whilst increasing the reversal amount requires a greater counter-trend price move to create a new column, and signal the end of a prevailing move, increasing the box size provides a coarser filter of all price movements. For example, Figure 5.6 shows a point and figure chart for the same data, but with a larger box size of \$0.04, and a 3-box reversal. In this case, only six columns are plotted on the chart. For completeness, Figure 5.7 presents the same specification of chart, with a \$0.04 box size, but with a 5-box reversal. Taking all of these examples together illustrates that there is a clear trade off between isolating important price movements and removing too much useful information.

Due to the advantageous way in which important price moves, trends, support and resistance levels are highlighted, the raw depiction of prices in a point and figure chart is useful. However, their value can be further enhanced by employing




Figure 5.6: Point and figure chart of stock price shown in Figure 5.3 with a box size of \$0.04 and reversal of 3 boxes

trading signals. The construction method allows buy and sell rules, based upon particular formations of columns of 'X's or 'O's, to be consistently employed. These patterns are also easier to identify in terms of a computer algorithm than patterns on standard price charts, such as the head and shoulders. This is because the point and figure has already been employed to identify 'significant' price moves and areas on the chart; thus, it is not necessary to use techniques such as kernel regression to identify maxima and minima. Therefore, in studying point and figure trading signals, it is possible to work on exactly the same basis as virtually all analysts using point and figure charts.

There are a number of different types of trading signals. At the most basic level, the initiation of a new column can be viewed as a buy or sell indication. For instance, the start of a new column of 'O's is seen as a break in an existing uptrend and the start of a new downtrend. Accordingly, this forms a sell signal. Conversely,




Figure 5.7: Point and figure chart of stock price shown in Figure 5.3 with a box size of \$0.04 and reversal of 5 boxes

the start of a new 'X' column signals a buy (a change from a downtrend to an uptrend).

More importantly, and given far greater attention by technical analysts, point and figure charts also provide buy or sell signals through patterns produced by the characteristics of the 'X' and 'O' columns on the charts. These patterns are expressed well by Zieg and Kaufman (1975), and adopted by Anderson and Faff (2008) in a study of point and figure trading rules and the S&P 500 futures contracts. This paper evaluates such patterns, which are specified in more detail below.

5.2.2 Point and figure literature

Over the past two decades, an increasing amount of research into aspects of technical analysis has been undertaken. Although previous studies have addressed topics such as moving averages, oscillators and price patterns, there has been very

little academic evaluation of the profitability of a trading strategy based on point and figure. Anderson and Faff (2008, p.2) note that “relevant literature on Point and Figure is extremely small—to our knowledge only three academic works have been published.” They recognise that two of these works are in German, by Hauschild and Winkelmann (1985) and Stottner (1990).

Hauschild and Winkelmann (1985) use daily data from 1970-1980 on 40 German securities. This study examines point and figure charts with a variety of box sizes, and also investigates five trading signals (corresponding to the five most common point and figure chart patterns). The results show that point and figure based trading signals are profitable. However, profitability is conditional on particular sub-periods. Specifically, it seems that buying or selling in response to signals is particularly rewarding when the market is in a trading range. In other circumstances, a simple buy-and-hold strategy is superior. Given that investors cannot know the state of the market in advance, this means that the authors conclude that point and figure has little value.

Whilst the results of this study are valuable, in the sense that they shed light on a previously scarcely touched upon area in the literature, there are several important limitations. First, the small sample size of 40 firms prevents a broader picture of the profitability of point and figure being obtained. It also means that the conclusions are weakened given that the number of trading signals produced is relatively low given the small number of securities investigated. Second, in using daily data, the study ignores the original and continuing application of point and figure to intraday data. Third, from the point of view of technical analysis, markets are very different for traders contemporaneously than in the 1970s. Traders now have an abundance of software for technical analysis, including plotting point and figure charts. It is therefore interesting to investigate point and figure in the current trading environment.

Using a larger sample of 445 German and foreign securities, Stottner (1990) uses

up to 18 years of data (depending on availability).² Their data is collected monthly; this is a significant disadvantage. First, because as already established, point and figure is intended for shorter time horizons. This does not accord with the much increased data availability for traders and investors, who can easily plot point and figure charts intraday or daily. Second, Stottner notes that the data contains only information on monthly highs and lows. This means that it is possible that more than one reversal has taken place within a month, yet by only having the single highest and single lowest price, it is impossible to determine. Therefore, any trading strategy may execute trades that would not be adopted by traders and/or that would be recorded later than would otherwise be the case with a greater frequency of observations. Furthermore, data on foreign securities is dominated in Deutschemarks; traders plotting point and figure charts are more likely to use the native currency.

Stottner compares the results of a point and figure strategy with a simple buy and hold strategy. The point and figure strategy is to trade when a reversal takes place. Thus, when a new column of 'X's ('O's) is started, when the current column is 'O's ('X's), a buy (sell) trade is entered. This is held until the signal reverses. Under the algorithm employed, this is approximated by employing a filter, which is selected to be 10%. In doing so, this is not strictly following the point and figure methodology discussed above. Furthermore, the choice of a 10% filter size is subjective. Perhaps most importantly, the study only looks at reversals in point and figure columns. This does not take into account well-established trading rules, and thus ignores key concepts such as support and resistance that can be clearly conveyed by point and figure charts. The results presented contrast the results from the point and figure strategy with simple buy and hold. It is found that point and figure does not outperform. Despite being one of a very small number of studies attempting to investigate point and figure, the drawbacks render the

²Samples from individual securities range from 10 to 18 years.

usefulness of this result doubtful.

There is a crucial disadvantage to both of the above studies. Point and figure was initially intended to be used on high-frequency data by floor traders who would form hand-drawn charts based on up-ticks and down-ticks throughout the trading day. However, in more recent times, the practitioner literature does demonstrate the use of point and figure with closing price data (Du Plessis, 2005, for example). The disadvantage of using end-of-day data is that information is lost about price changes during the trading day. Even if we have data on opening, high, low and closing prices, it is unknown whether the high occurred before the low (or vice versa) and thus whether a reversal occurred within the day.

Attracted by the useful filtering approach, there have been academic applications of point and figure that, whilst using an element of the technique, do not seek to evaluate a trading strategy. Elliott and Hinz (2002) make use of point and figure to investigate portfolio optimisation. They use point and figure as a method of identifying “significant times” where re-balancing of a portfolio is required. This indicates the benefit of the use of point and figure as a method of deriving times of analytically important price change from noisy data. Giles (2005) applies point and figure charting to monetary policy in terms of forecasting UK interest rates. Giles determines that the reversals seen on point and figures charts have a unique appeal in terms of filtering. Likewise, the breakouts from the various point and figure patterns are seen to be valuable. Whilst both of these studies do not view point and figure charting from the point of view of a trading strategy, they are important in underlining the unique approach of point and figure, and the value of its approach to filtering noisy data.

There may be several reasons for this. First, as illustrated by Figure 5.1 as opposed to Figure 5.2, point and figure charts have a drastically different visual appearance from standard line, bar and candlestick charts. Second, point and figure charts are less easily evaluated programmatically. For example, it is relatively trivial

to construct a program to derive buy and sell signals from a simple moving average crossover. Indeed, such signals can be derived from simply evaluating a vector of prices for a given security.

Finally, point and figure charting was originally designed to be a reliable and informative method for floor traders to track price movements by hand. Accordingly, point and figure chartists worked from intraday data. As will be seen below, the limited point and figure literature does not make use of real-time data. This is understandable, as it is only recently that computational power has been sufficient to evaluate the millions of observations from trade databases. However, it leaves a significant gap in our understanding, which is addressed by the empirical work in this study.

Anderson and Faff (2008) provide the most recent investigation of point and figure. Their study goes some way towards handling some of the problems with previous work that were identified above. Using 1-minute data for S&P 500 futures contracts traded between 1990 and 1998 allows for an initial analysis of point and figure with the high frequency data for which it was originally intended. Furthermore, instead of merely using point and figure in the context of a filter, as implemented by Stottner (1990), they evaluate a number of trading strategies based upon point and figure chart patterns.

Their study refers to these patterns as trading rules, and the specification for each of these is adopted from Zieg and Kaufman (1975), an important practitioner text. This is a sensible approach, as Zieg and Kaufman aimed their text at traders, and it provides clear information that these rules were available to traders over a long period of time. Anderson and Faff record mixed results; whilst there is some limited evidence of profitability, this was not a uniform result across all years in the sample. Further, profits that were shown appear to owe to periods of trading with relatively high volatility.

However, while Anderson and Faff clearly advance our knowledge of point

and figure, there are some important limitations. First, in looking solely at the S&P 500 futures, we still lack any knowledge of the profitability of point and figure trading strategies for individual securities. This is important as traders apply the technique to stocks. Furthermore, whilst the use of 1-minute data allows an initial investigate into the profitability of point and figure at a much higher frequency than previous work, it is still unknown how well the rules perform with ‘real-time’ data. A further point, discussed below, is that with 1-minute data we do not know how price moves *within* the 1-minute periods. This is important, as if a particular price in the intervening period would have triggered a new plot on the chart, but this move is reversed before the next observation, we have no knowledge of it. Using all data points negates this problem.

To address the above points, this study takes the approach of using a large sample of individual securities—the constituents of the S&P 500—paired with un-aggregated data from the NYSE TAQ database.³ By using this ultra high-frequency data it is possible to extend our understanding of this area of technical analysis. Whilst being of interest in the context of market efficiency, this study makes an even wider contribution, in terms of being of interest to traders who use technical analysis to make intraday investment decisions.

5.3 Data and methodology

NYSE trade data for S&P 500 stocks was used from January 1 to December 31, 2005, with data downloaded from the consolidated trades database. This database contains all data on trades for stocks on the NYSE. The list of S&P 500 constituents was re-evaluated monthly in order to take into account additions and deletions from the index. Observations were collected from 9:30 a.m. EST to 4:00 p.m. EST. ‘Late’ trades that are reported to the tape some time after actually occurring (a

³Although the data is un-aggregated (i.e. is used trade-by-trade rather than being filtered to 1-minute or 5-minute observations), ‘cleaning’ takes place to eliminate data errors. The methodology for this is discussed below.

Sale Condition field of 'Z') are removed from the sample. Similarly, trades that occurred in sequence but are reported at a later time (a Sale Condition field of 'O') are also excluded. Only trades where the Correction Indicator is equal to 0 or 1—regular trades and trades where data is subsequently corrected, respectively—are used, thus removing pre-identified errors from the data. These steps mitigate the inherent problems with using the NYSE ultra-high-frequency trades database (Brownlees and Gallo (2006) identify the issues that result from not employing such data cleaning).

As detailed above, point and figure charts are conditioned on two variables: the box size, B , and the reversal amount, R . The larger the box size, the greater the filtering impact of constructing a point and figure chart, i.e. if there is an inverse relationship between B and the number of observations. This can be seen visually by looking back to Figures 5.4 and 5.6. These use intraday Microsoft price data to illustrate the reduction in the number of plot points, when increasing box size from \$0.02 to \$0.04. Intuitively, a greater reversal amount implies that a larger counter-trend price movement is needed before a new column is plotted; this also implies an inverse relationship between R and the number of resulting observations point and figure chart plots. Figures 5.4 and 5.5 visually display the effects of increasing the value of R from 3 to 5 boxes on a sample of intraday stock data.

The first step in the analysis is identifying columns of 'X's and 'O's to develop a point and figure chart. First, R and B are specified. Data is read from the trades database, and at data point, P_t , where a price change occurs (recall that times when prices do not change are ignored in point and figure analysis), and a number of steps are performed. If the current column, C_j , is 'X's and the price P_t has increased if $P_t \geq (C_j^{high} + B)$, then $P_t - C_j^{high}$ is evaluated and rounded down to the nearest box size, which then becomes the new C_j^{high} . The procedure for a price fall when the column is 'O's is analogous, but compares P_t to C_j^{low} .

Where the current column is 'X's and the price has fallen, if $P_t > [C_j^{high} -$

$(R \times B)]$ then the price move is not greater than the reversal amount so nothing is recorded. However, if $P_t < [C_j^{high} - (R \times B)]$ then a new column of 'O's is begun (C_{j+1}). Again, the process is analogous when the current column is 'O', apart from that we compare P_t with $C_j^{high} - (R \times B)$ to see if a reversal is recorded, leading to a new column of 'X's.

Once the columns have been identified, together with their highs and lows, it is possible to apply algorithms to compute the buy and sell returns of applicable trading rules. We choose the four most common trading rules that are detailed in the practitioner literature.

Zieg and Kaufman (1975) develops a set of point and figure patterns which can be interpreted as trading rules. Generally, these rules are developed around breakouts from levels set by preceding columns of 'Xs' and 'Os'. These patterns are clearly recognised by the practitioner literature (Dorsey, 2007, for example) and their longstanding use goes some way to mitigating concerns over data mining. Anderson and Faff (2008) also adopts these definitions in their investigation of point and figure charting and the S&P 500 futures contract.




Figure 5.8: Double Top buy signal (B1)

The least restrictive patterns are the Double Top and Double Bottom formations, illustrated in Figure 5.8 and Figure 5.9, respectively. Dorsey (2007) makes the connection between key levels on point and figure charts and important support and resistance areas. In the Double Bottom, three columns on the chart are required.




Figure 5.9: Double Bottom sell signal (S1)

The first and last columns are 'Xs', with the second 'Os'. The Double Top buy signal is given when an 'X' is recorded in the third column above the level of the highest 'X' in the first column. Technical analysts' intuition behind this pattern is that when price advances beyond the level of resistance, previously encountered at the high of the first column, that a breakout has begun. Similarly, the Double Bottom records a fall through a previous level of support shown by the first column of 'Os' and a sell signal is given. These formations are labelled B1 and S1, following Zieg and Kaufman (1975) and Anderson and Faff (2008).

We now add the condition of a rising bottom to the Double Top and a declining top to the Double Bottom. The practitioner literature, referenced above, determines that the Double Top with rising bottom pattern provides a stronger signal to traders because the rising bottom indicates that supply pressure is becoming more easily overcome in recent price moves. For the signal to be triggered, demand overwhelms previous levels (shown by the previous highs on 'Xs' columns). This pattern is exhibited in Figure 5.10. For the Double Bottom with declining top, supply side pressure on prices is increasing, and is finally overcome as price breaks below the previous low of the penultimate 'O's' column, triggering a sell signal as shown in Figure 5.11.

As the name suggests, the Triple Top and Triple Bottom imply that price has reached a particular level on the charts for a third time, and then broken it, to

			X	← Buy
	X		X	
O	X	O	X	
O	X	O	X	
O	X	O		
O				

Figure 5.10: Double Top buy signal with rising bottom (B2)

X				
X	O	X		
X	O	X	O	
	O	X	O	
	O		O	
			O	← Sell

Figure 5.11: Double Bottom sell signal with declining bottom (S2)

				X	← Buy
X		X		X	
X	O	X	O	X	
X	O	X	O	X	
	O		O		

Figure 5.12: Breakout of triple top (B3)

	X		X		
O	X	O	X	O	
O	X	O	X	O	
O		O		O	
				O	← Sell

Figure 5.13: Breakout of triple bottom (S3)

produce a buy or sell signal. For the Triple Bottom, shown in Figure 5.12, price has overcome a level of resistance marked by the two previous columns of 'X's. Price has returned to a level of support three times in the Triple Top, shown in Figure 5.13, and broken through this level to generate a sell signal.

These two patterns are extended in the Ascending Triple Top and Descending Triple Bottom shown in Figure 5.14 and Figure 5.15, respectively. The Ascending Triple Top requires the lows of successive 'O's columns to be higher, whereas the Descending Triple Bottom requires the highs of successive 'X's columns to be lower. In a similar fashion to the Ascending Double Top and Descending Double Bottom, these two patterns present a stronger signal to traders, who perceive that the balance between supply and demand influencing prices is convincingly changing.

As detailed above, the box size (B) and reversal amount (R) are key features

				X	← Buy
		X		X	
X		X	O	X	
X	O	X	O	X	
X	O	X	O		
	O				

Figure 5.14: Ascending Triple Top (B4)

	X				
O	X	O	X		
O	X	O	X	O	
O		O	X	O	
		O		O	
				O	← Sell

Figure 5.15: Descending Triple Bottom (S4)

of the point and figure technique. Together, they control how harshly the process filters data; a large box size will serve to capture only the major moves, and a large reversal amount requires a larger counter-trend movement for a new column to begin. Given the significance of these two variables, it is helpful in our understanding of point and figure to investigate the impact of adjusting the values of B and R .

Initially, the box size is set at one cent ($B = 1¢$). This is reasonable given that this study looks at point and figure charting with high frequency trade data—as of January 29, 2001, the NYSE moved to full decimalisation and the minimum tick size became one cent. Box sizes of two cents ($B = 2¢$) and four cents ($B = 4¢$) are also evaluated. It is to be expected that as box size is increased, fewer columns of ‘Xs’ and ‘Os’ will be plotted on the point and figure chart. Thus we might expect that (relatively) more important price moves will be captured.

Similarly, two different reversal amounts are investigated. First, the traditional three-box method ($R = 3$). This requires a reversal equivalent to three boxes for a new column to be initiated. Second, the trading rules are applied to a five-box construction method ($R = 5$). A price move equivalent to five boxes is required for a new column. It is also expected that increasing the reversal amount will increase the ‘coarseness’ of the point and figure filter. It is important to investigate this aspect as, whilst there is a clear consensus in the practitioner literature regarding trading signals, there is less consensus about the value of R . However, the three and five box construction methods are by far the most common.

After constructing point and figure charts, trading rules are evaluated for profitability. In accordance with the practitioner literature (Zieg and Kaufman, 1975; Dorsey, 2007, for example), these rules are treated symmetrically. Therefore, when a buy signal is generated by a Double Top (B1), the position is held until the occurrence of a Double Bottom (S1). Returns are then computed; recall that the entry and exit points reflect the actual price that caused the relevant box to be filled

to complete the pattern (and not the box prices themselves). Thus, upward bias in profitability is avoided in the results. Returns for a particular trading signal, i , are continuously compounded. For example, the return for buy signal $B1$ is computed over the ensuing period, until an $S1$ sell signal is identified, as

$$r = \ln P_{t=S1} - \ln P_{t=B1} \quad (5.1)$$

As a $B1$ signal initiates a short sale, to make the findings clearer we reverse the formula so that, for sells, a successful trade is shown as a positive return. Using the case above, the return for selling short after an $S1$ signal and covering when a $B1$ signal is observed is

$$r = \ln P_{t=B1} - \ln P_{t=S1} \quad (5.2)$$

As noted in Anderson and Faff (2008), building upon the practitioner point and figure literature, it should be noted that the trading rule patterns encompass each other. For example, by definition, the Triple Top ($B3$) contains a Double Top ($B1$). This can be seen by comparing Figure 5.8 and Figure 5.12. Therefore, multiple long or short positions may be initiated at the same price as several rules may be triggered concurrently.

Anderson and Faff (2008) correctly point out that slippage is important in constructing and interpreting the results from a point and figure strategy. Slippage occurs when a trading signal is generated by a point and figure pattern, but price has already advanced beyond the prices that bound a particular box on the chart. For example, if price is required to reach \$50 for a buy signal to be triggered, but the trade that actually surpasses this level for the first time occurs at \$50.10, then if we were to record the buy point at \$50, profitability would be overstated. Anderson and Faff (2008) instead take the actual price that triggered the move to the next box as the buy price (\$50.10 in the above example). Whilst their research deals

with one-minute data for the S&P 500 futures contract, this study deals with ultra-high-frequency trades data. However, although likely to be smaller in magnitude, slippage may still be an issue. Accordingly, the buy price is similarly determined to be the actual price on the tape that triggered the filling of a point and figure box, causing a buy or sell signal to be triggered.

In this study, we make a major contribution in dealing with tick-by-tick data—the original preserve of point and figure charting. Accordingly, all round-turn trades occur within the same trading day; positions are not carried overnight.

5.4 The Profitability of Point and Figure Trading Rules

This section presents the returns of the point and figure trading strategies detailed in the previous section, using data from the NYSE consolidated trades database for 2005, for all S&P 500 constituent stocks. To allow clearer analysis of profitability, results are segregated in two ways based on the construction of the point and figure charts that gave rise to trading signals. First, two reversal values are tested. $R = 3$ corresponds to a three-box reversal point and figure chart, and $R = 5$ a five-box chart. Second, for each of $R = 3$ and $R = 5$, results for three different point and figure box sizes are evaluated. These are 1¢ , 2¢ and 4¢ , which are denoted as $B = 1$, $B = 2$ and $B = 4$, respectively. Within each of these scenarios, results are shown for the four major trading rules under investigation: the Double Top and Double Bottom, the Double Top with rising bottom and the Double Bottom with declining bottom, the Triple Top and the Triple Bottom, and the Ascending Triple Top and Descending Triple Bottom. Presenting results in this manner allows scrutiny of the sensitivity of point and figure analysis to the choice of reversal amount and box size. This is of interest because, as detailed above, a greater reversal amount and box size leads to increased filtering of price data.

5.4.1 3-box chart construction

Results for the three-box point and figure chart construction method, $R = 3$, are presented in Table 5.1. Corresponding to a three-box point and figure chart, the reversal value of $R = 3$ means that price must 'reverse' by three times the box size, B for a new column to be initiated. The columns show the three different box sizes of $B = 1\text{¢}$, $B = 2\text{¢}$ and $B = 4\text{¢}$. The rows of the table are broken down into groups representing the four point and figure trading rules that are evaluated. For each rule, the number of buys and sells generated are shown as N Buy and N Sell, respectively. The mean buy and sell returns are presented together with t -statistics and corresponding p -values ($Pr > |t|$). % of prof buys and % of prof sells represent the fraction of profitable trades for the buy and sell rules, respectively.

The B1 and S1 rule buys when there is a breakout from a Double Top and sells when there is a breakout from a Double Bottom. A very large number of buy and sell trades are generated over 2005, with 1,629,912 buys on breakouts of Double Tops and 1,626,084 sells on breakouts from Double Bottoms. There are therefore approximately 13 buy and sell signals given per S&P 500 security, per day. Given that the B1 and S1 rules are the least restrictive, the high number of trading signals is unsurprising; this concurs with results presented by Anderson and Faff (2008), who also found far more trades for B1 and S1 rules than for the more restrictive rules. The similarity between the number of buy and sell trades serves to confirm the rationale of a symmetrical trading strategy, where buys from Double Top breakouts are sold on a corresponding Double Bottom breakout.

The results for the B1 and S1 rule show that both the mean buy and mean sell returns are negative, albeit the mean sell return is slightly smaller in magnitude than the buy return. This shows that both buys and sells entered into according to this rule generate losses overall. Both the mean buy and mean sell returns are significantly different from zero. Only around a third of the buy and sell trades were successful; 33.9% of buy and sell trades were profitable. However, traders

Table 5.1: Returns of Point and Figure Trading Rules, where $R = 3$

Strategy	$R = 3, B = 1¢$	$R = 3, B = 2¢$	$R = 3, B = 4¢$
<i>B1 & S1 rule</i>			
N Buy	1629912	600038	168487
Mean buy rtn.	-0.01245	-0.04527	-0.12402
t Value	-60.60	-92.94	-92.46
$Pr > t $	<.0001	<.0001	<.0001
% prof buys	33.9	31.7	28.7
N Sell	1626084	600147	169581
Mean sell rtn.	-0.01026	-0.03999	-0.10727
t Value	-48.94	-81.92	-80.96
$Pr > t $	<.0001	<.0001	<.0001
% prof sells	33.9	32.3	29.9
<i>B2 & S2 rule</i>			
N Buy	387168	130660	32807
Mean buy rtn.	0.06741	0.06936	0.04929
t Value	125.85	54.21	14.06
$Pr > t $	<.0001	<.0001	<.0001
% prof buys	47.3	45.0	41.3
N Sell	389934	131662	33315
Mean sell rtn.	0.07069	0.07603	0.06118
t Value	133.50	61.66	17.92
$Pr > t $	<.0001	<.0001	<.0001
% prof sells	47.5	45.9	42.6
<i>B3 & S3 rule</i>			
N Buy	256876	83870	21572
Mean buy rtn.	0.00730	-0.04339	-0.17467
t Value	7.27	-19.40	-31.97
$Pr > t $	<.0001	<.0001	<.0001
% prof buys	39.3	35.8	28.4
N Sell	257422	84518	21312
Mean sell rtn.	0.01444	-0.02161	-0.13581
t Value	14.06	-9.45	-25.01
$Pr > t $	<.0001	<.0001	<.0001
% prof sells	39.6	36.6	30.8
<i>B4 & S4 rule</i>			
N Buy	39443	8288	1049
Mean buy rtn.	0.05934	0.03296	-0.00998
t Value	28.98	5.42	-0.42
$Pr > t $	<.0001	<.0001	0.6760
% prof buys	43.9	40.2	36.0
N Sell	39732	8389	1051
Mean sell rtn.	0.06446	0.04987	0.02921
t Value	31.30	7.97	1.23
$Pr > t $	<.0001	<.0001	0.2174
% prof sells	44.3	40.6	37.0

The sample period is January 1 2005 to December 31 2005, comprising all of the constituents of the S&P 500, with 'real-time' from the consolidated trades database. The trading rules (e.g. B1 and S1) correspond to those discussed in the text. *N Buy* (*N Sell*) is the number of buy (sell) trades over the year. $Pr > |t|$ is the *p*-value. % *prof buys* and % *prof sells* represent the percentage of profitable buy and sell trades entered into according to the rule, respectively. This table shows results from point and figure charts constructed with a 3-box technique ($R = 3$). Results are shown for three box sizes ($B = 1¢$, $B = 2¢$ and $B = 4¢$). Mean buy and sell returns are multiplied by 100 for ease of interpretation. Thus, the first mean buy return is -0.1245%.

would not be able to profitably use the B1 and S1 rule to make investment decisions.

Given that the B1 and S1 rule is the least restrictive it may not be entirely surprising that profits are not exhibited from Double Top buys and Double Bottom sells. The B2 and S2 rule imposes a further constraint: the Double Top should have a rising bottom and the Double Bottom should have a declining top. Unsurprisingly, this constraint substantially reduces the number of trades reported, to about a quarter of those generated by the B1 and S1 rule. 387,168 buys and 389,934 sells are shown, a mean average of just over three buy and three sell signals per S&P 500 constituent, per trading day. Again, the similar number of buys and sells lends support to the symmetrical strategy that is employed.

Unlike the B1 and S1 rule, the returns for sells and buys are both positive; they are also substantially greater in size. The Double Top with rising bottoms breakout produced a mean return of 0.06741% per trade, and Double Bottom with declining top breakouts produced a mean return of 0.07069% per trade. Both results are significantly different from zero. Imposing the additional restriction produces more trades that are profitable: 47.3% of buy trades and 47.5% of sell trades. These results suggest that this trading rule captures potentially useful information.

The B3 and S3 rule also shows mean buy and sell profits. Imposing further restrictions to analyse breakouts from Triple Tops and Triple Bottoms reduces the number of trades from both buys and sells; just over two buy and two sell trades per day for each security are recorded, on average. Given the importance attached to point and figure patterns as signals to buy and sell in the practitioner literature, it is interesting to see whether imposing greater restrictions to detect more complex patterns increases profitability. The percentage of profitable buy and sell trades from the B3/S3 rules has fallen compared to the B2/S2 rules suggesting that, in this case, the further constraints have not improved the ability to identify profitable trades. This is confirmed by smaller mean buy and sell returns, of 0.00730% and 0.01444% per trade, respectively. This result shows that it is not a

given that imposing further constraints on the pattern formations, and detecting more complex patterns, increases mean returns.

The B4 and S4 rule generates buys and sells from breakouts of Ascending Triple Tops and Descending Triple Bottoms (similar to the progression from the B1 and S1 to the B2 and S2 rule). The number of trades recorded is again reduced to 39,443 and 39,732 for buys and sells, respectively. This corresponds to approximately one buy and one sell trade every three days, per security. Given the more complex specification of the Ascending Triple Tops and Descending Triple Bottoms, this is unsurprising. The results are interesting in the sense that the mean buy and sell return per trade is not as great as under the less restrictive specification of the Ascending Double Top and Descending Double Bottom. Furthermore, there has been a reduction in the fraction of profitable trades. Therefore, the greater selectivity, which translates into fewer trading opportunities, does not lead to greater profitability.

These results show that, with the exception of the least restrictive strategy, the trading signals produced by point and figure charts are profitable. This is based on a three-box reversal and a box size of 1¢ ($R = 3, B = 1\text{¢}$). As the box size is a crucial element of the point and figure technique, with an inverse relationship between box size and the number of elements on a point and figure chart (from which trading signals are derived), it is important to investigate the impact of a change in its size.

Increasing the box size from a penny to two pennies, $R = 3, B = 2\text{¢}$, unsurprisingly shows a fall in the number of buy and sell signals found. This is expected as the point and figure filter becomes 'coarser' with the increase in box size. Under the B1 and S1 rule, the number of buy and sell trades are just over a third of those for $B = 1\text{¢}$ (about 5 trades each way per day, per security, compared to about 13 for the smaller box size). The results show that, for the simplest case of the B1/S1 rule, increasing the box size from 1¢ to 2¢ actually causes a roughly four-fold rise

in the mean loss from each buy and sell trade. The same picture is seen when the box size is again increased to 4¢ . The average number of buy and sell signals given per day falls further to around 1.4, and there is a relatively big increase in the mean loss per trade. In fact, compared to $B = 2\text{¢}$ where the mean buy (sell) return was -0.04527% (-0.0399%), $B = 4\text{¢}$ shows a much larger mean buy (sell) return of -0.12402% (-0.10727%). However, this result may be explained by there being fewer buy and sell signals. As trades are exited on an opposing signal (i.e. when a B1 pattern causes trade entry, an S1 pattern causes trade exit), this means that the average holding period is similarly increased. However, as with previous box sizes, this pair of rules is not profitable.

For the B2 and S2 rule, increasing box size from $B = 1\text{¢}$ to $B = 2\text{¢}$ and, subsequently, to $B = 4\text{¢}$, causes a similar reduction in the average number of trades per security, per day. However, whilst the increase from $B = 1\text{¢}$ to $B = 2\text{¢}$ marginally increases both the mean buy and mean sell returns, there is actually a reduction in profitability when moving to $B = 4\text{¢}$. There are two possible reasons for this. First, making the filtering element of point and figure more restrictive by increasing box size may have removed some profitable trading signals. Second, as trades are closed when an opposing signal is recorded, the marked reduction in mean trades per security per day translates into a longer mean trade duration. The implication is, therefore, that the usefulness of point and figure trading signals may decay relatively quickly.

A different and interesting picture is presented for both the B3 and S3 rule when box size is increased from $B = 1\text{¢}$ to $B = 2\text{¢}$ and $B = 4\text{¢}$. Again, the number of trading signals falls markedly as box size is increased. However, when box size is increased, the mean buy and mean sell returns per trade now become negative. Whilst the B3 and S3 rule was profitable for the smallest box size under investigation, when this is increased, it is no longer the case. Indeed, when $B = 4\text{¢}$, the mean loss per trade for both buys and sells is comparatively very large in

magnitude. This result is supported by a relatively lower fraction of profitable buy and sell trades of 28.4% and 30.8%, respectively. Again, it is likely that the increased length of time between entry and exit of trades is a factor. This result shows that the choice of box size is extremely important to point and figure profitability.

The results for increased box size for the most restrictive rules, B4 and S4 (breakouts from Ascending Triple Tops and Descending Triple Bottoms), do not concur, however. The mean buy and mean sell returns were profitable under box size $B = 1\zeta$. When box size is increased to $B = 2\zeta$, both buys and sells are still profitable, although the mean returns for each is reduced. This corresponds with a similar proportional reduction in the fraction of profitable trades. With a box size of $B = 4\zeta$ for the B4/S4 rule, the mean sell trade return is still positive but the mean return for buys becomes negative. However, these are the only mean returns presented that are insignificant. This is partly due to there only being 1,049 and 1,051 buy and sell trades recorded, respectively. It would therefore be unwise to place emphasis on this particular result. However, the small number of trades does show that the choice of box size is a crucial factor.

Taken together, the results for point and figure trading strategies where charts are constructed with a three-box method, $R = 3$, allow some interesting overall observations to be made. It is shown that point and figure trading strategies based upon the four rules under investigation are profitable in the majority of cases. As expected, both increasing the constraints of the trading rules to focus on more complex patterns, and increasing the box size, has a pronounced effect on the number of signals generated. The choice of trading rule is also important; for all box sizes, the simplest B1/S1 rule was not profitable. The most consistently successful strategy is based on the B2/S2 rule. It is evident that these results do not support the lack of point and figure profitability shown in the limited previous research. The findings support some point and figure trading rules being able to successfully identify important areas of support and resistance.

5.4.2 5-box chart construction

The results presented in this section address the issue of whether the reversal amount chosen in the construction of point and figure charts influences the profitability of point and figure trading strategies. To examine this, having already discussed results for the three-box construction method ($R = 3$), Table 5.2 presents findings for the five-box point and figure chart ($R = 5$). In this case, price is required to reverse by a greater amount (five versus three times box size, B) for a new column to be initiated. A larger reversal amount increases the degree of filtering of price data (this can be seen visually by referring to the charts presented earlier, for example Figure 5.4 as opposed to Figure 5.5).

It is shown that the number of signals generated from the five-box construction method, $R = 5$, are much lower than for $R = 3$. Taking a box size of $B = 2\text{¢}$ as an example, there are 924,644 buy trades for the B1/S1 trading rule combination with $R = 5$ in contrast to 1,629,912 for $R = 3$. This equates to a reduction in the average number of buy signals per day from approximately 13 to 7 using a three and five-box reversal chart, respectively. As such, it is interesting to establish if this has filtered out some of the less profitable trades.

For the least restrictive B1 and S1 strategy, the fraction of profitable buy and sell trades has uniformly decreased for all three box sizes evaluated. As with $R = 3$, all of the mean buy and sell returns are negative. However, they are substantially greater in magnitude under the five-box construction method. For instance, the mean buy return per trade for $R = 5, B = 4\text{¢}$ is -0.22572% compared to -0.12402% reported for $R = 3, B = 4\text{¢}$. The B1/S1 strategy leads to mean losses for both buys and sells across all box sizes, and is invariant to the size of the reversal amount used for point and figure chart construction.

As with the three-box chart results, the mean return per trade for both buys and sells under the B2/S2 rule is positive in all cases. Again, due to the coarser filter imposed by a greater reversal amount, the number of reported trades is

Table 5.2: Returns of Point and Figure Trading Rules, where $R = 5$.

Strategy	$R = 5, B = 1¢$	$R = 5, B = 2¢$	$R = 5, B = 4¢$
<i>B1 & S1 rule</i>			
N Buy	924644	290330	64907
Mean buy rtn.	-0.03104	-0.08680	-0.22572
t Value	-91.84	-98.46	-81.71
$Pr > t $	<.0001	<.0001	<.0001
% prof buys	32.0	29.5	25.3
N Sell	923042	291281	65352
Mean sell rtn.	-0.02724	-0.07574	-0.19791
t Value	-79.05	-86.47	-73.99
$Pr > t $	<.0001	<.0001	<.0001
% prof sells	32.4	30.4	27.0
<i>B2 & S2 rule</i>			
N Buy	225934	64743	12623
Mean buy rtn.	0.07542	0.07347	0.02692
t Value	90.92	33.35	3.78
$Pr > t $	<.0001	<.0001	0.0002
% prof buys	46.7	44.1	39.3
N Sell	227493	65927	12876
Mean sell rtn.	0.08030	0.08175	0.05451
t Value	96.98	39.21	8.21
$Pr > t $	<.0001	<.0001	<.0001
% prof sells	47.5	45.1	42.0
<i>B3 & S3 rule</i>			
N Buy	72252	18740	3430
Mean buy rtn.	-0.02439	-0.00114	-0.35709
t Value	-9.20	-18.70	-22.18
$Pr > t $	<.0001	<.0001	<.0001
% prof buys	38.0	31.4	22.8
N Sell	73434	18720	3354
Mean sell rtn.	-0.00196	-0.06245	-0.29049
t Value	-0.78	-10.25	-16.55
$Pr > t $	0.4349	<.0001	<.0001
% prof sells	38.6	34.0	25.2
<i>B4 & S4 rule</i>			
N Buy	16106	2853	13
Mean buy rtn.	0.05642	0.01312	-0.17564
t Value	14.88	1.02	-0.48
$Pr > t $	<.0001	0.3086	0.6368
% prof buys	42.3	36.6	69.2
N Sell	16136	2818	17
Mean sell rtn.	0.07254	0.03606	-0.24094
t Value	18.66	2.84	-0.65
$Pr > t $	<.0001	0.0046	0.5228
% prof sells	43.8	40.0	35.3

The sample period is January 1 2005 to December 31 2005, comprising all of the constituents of the S&P 500, with 'real-time' from the consolidated trades database. The trading rules (e.g. B1 and S1) correspond to those discussed in the text. *N Buy* (*N Sell*) is the number of buy (sell) trades over the year. $Pr > |t|$ is the *p*-value. % *prof buys* and % *prof sells* represent the percentage of profitable buy and sell trades entered into according to the rule, respectively. This table shows results from point and figure charts constructed with a 5-box technique ($R = 5$). Results are shown for three box sizes ($B = 1¢$, $B = 2¢$ and $B = 4¢$). Mean buy and sell returns are multiplied by 100 for ease of interpretation. Thus, the first mean buy return is -0.03104%.

considerably lower. The fraction of profitable buy and sell trades has decreased for all box sizes (with the exception of sell trades for $B = 1\text{¢}$, where the fraction is the same). Interestingly, although the ratio of successful to unsuccessful trades has fallen, the mean buy and sell returns have increased for both $B = 1\text{¢}$ and $B = 2\text{¢}$. This suggests that whilst changing from a three to five-box reversal method identifies slightly fewer profitable trades (as a proportion of the total), the mean returns from those that are profitable have increased. Yet, this is not the case for $B = 4\text{¢}$, where mean buy and sell returns per trade have decreased. In particular, the mean buy return under the B2/S2 strategy for $R = 5, B = 4\text{¢}$ is 0.02692% compared to 0.04929% for $R = 3, B = 4\text{¢}$.

The B3 and S3 strategy proved to be the least successful for a three-box chart construction method. This is still the case for five-box charts, but the result is even more pronounced. In this case, all of the mean buy and sell returns are negative, although the mean sell return for $B = 1\text{¢}$ is insignificantly different from zero. Furthermore, the fraction of profitable trades is relatively small. For instance, only 22.18% of buy trades for $B = 4\text{¢}$ are profitable.

Breakouts from Ascending Triple Tops and Descending Triple Bottoms are profitable for box sizes of $B = 1\text{¢}$ and $B = 2\text{¢}$. The mean returns for $B = 4\text{¢}$ are not open for interpretation, given that only 13 buy and 17 sell trades were reported over the period. However, this result does serve to demonstrate that the reversal amount has a large impact on the number of trading signals generated, and should be regarded as important by traders. For $B = 1\text{¢}$ and $B = 2\text{¢}$, mean buy return per trade is lower than under the three-box method. Mean sell returns are higher for $R = 5$ than $R = 3$ for a $B = 1\text{¢}$. The picture for the B4/S4 strategy is therefore somewhat mixed; whilst some returns are lower (and one significant return higher) under $R = 5$ versus $R = 3$, there are far fewer trading opportunities. However, this may be welcomed in the face of transactions costs, so the larger box size may still be preferred.

5.5 Conclusions

Point and figure is one of the oldest forms of technical analysis and is still important to traders today. This study investigates the point and figure technique using ultra-high-frequency intraday data for a sample of 500 large US stocks. This is in sharp contrast to the limited previous research in this area which does not look at intraday data, with the notable exception of Anderson and Faff (2008), although their study is restricted to one futures contract.

It is seen that point and figure works to filter price data as expected: increases in box size and a greater reversal amount employed in the construction technique reduce the number of trading signals recorded. The results show that profits are available to day traders in S&P 500 stocks. These profits are well represented across the differing box sizes and reversal amounts that are evaluated. It is seen that the least restrictive patterns of the Double Top and Double Bottom are consistently loss-making. However, the more restrictive rules show significant profits in most cases for the most commonly used (by traders) three-box chart. The most successful rule pair, B2 and S2, suggests that the point and figure technique is successful in isolating areas of support and resistance. Whilst the magnitude of returns is relatively small, it should be remembered that these are intraday trades. Traders with low transaction costs and liquidity traders could profitably employ point and figure methodologies.⁴

This is a valuable result for several reasons. First, as explained at the start of this study, our understanding of point and figure up to now has been relatively limited. Second, we know nothing at all about point and figure profitability in the form of its originally intended usage on 'real-time' data. Third, the point and figure technique has a history of over 100 years usage, and the trading rules tested

⁴In the previous chapter, which looked at intraday relative strength, it was noted that transaction costs are an important factor. However, short-term day trades who employ technical analysis trade very frequently, with survey evidence showing that a significant proportion do so profitably, net of transaction costs. Also, the costs of short selling may be relevant, and could be considered by future research.

here are well established, yet are still successful. As point and figure charting is a form of technical analysis and only utilises past price data, this result is therefore interesting in relation to market efficiency.

Chapter 6

Conclusions

The introduction to this thesis noted that academics have traditionally been highly sceptical of technical analysis. Thus, existing work is limited in many areas. The four distinct empirical chapters of this work make an extensive contribution to our understanding of technical trading strategies.

6.1 Thesis overview and importance of findings

The first empirical chapter investigated the head and shoulders, which is the most important and well-known chart pattern. The existing state of research was extended considerably in a number of important ways. First, the new concept of the trade lag was developed, which allowed investigation of the importance of the length of time between formation of head and shoulders patterns and the ability to identify them. Second, the value of a kernel smoothing approach, to identify local peaks and troughs in price data, was recognised. However, this was applied in a new fashion. Most importantly, by using locally optimised bandwidth, the subjective alteration of the degree of smoothing in previous work could be avoided. Third, by developing faster algorithms, it was possible to undertake bootstrap analysis. Fourth, a broad insight into the profitability of chart patterns was presented with the use of a large sample of data for individual UK stocks.

It was shown that the head and shoulders pattern provides economically valuable information. However, this is contingent on the time horizon for which trades are held, and how recently patterns are identified. Limitations in previous work have not allowed this distinction to be made. Head and shoulders tops were very successful; for example, annual excess returns of around 2% were found when holding trades for 60 days. The trade lag shows that the time period between detection and trading on patterns is vital. With this in place, the mean excess return becomes 3.5% on an annualised basis. Head and shoulders bottoms were not profitable at longer time horizons, but significant excess returns were present

for holding periods of less than 10 days.

Chapter 3 built upon and considerably extended the work in Chapter 2 through a number of significant innovations. Following analysis of the practitioner literature, two new specifications were developed for the head and shoulders pattern. According to technical analysts, head and shoulders patterns can signal either a continuation or reversal of an existing trend; this aspect has been ignored by academic research up to now. Findings showed that, contrary to technical analysts' beliefs, this is not important. Emphasis is also placed on the neckline, which is regarded as an important confirmatory aspect of head and shoulders patterns. When this was included in the pattern specification, it was found that mean excess returns increased. For example, buy signals from the inverse head and shoulders pattern generated mean excess returns of 5.5% on an annual basis. Furthermore, this chapter also evaluates a longer formation period for patterns, of up to 65 days. According to the practitioner literature, the longer the time period over which patterns form, the more important they are. The empirical results show this not to be the case. Indeed, mean excess returns from a shorter formation period of 35 days are greater. Bootstrapping demonstrated the significance of these findings, and also suggested that traders using the head and shoulders pattern do not appear to be subjecting themselves to increased risk, in order to receive excess returns. These findings are clearly contrary to weak-form market efficiency. Technical analysts are correct in asserting the importance of the head and shoulders pattern, but the findings of this study suggest it may not function in exactly the way they think.

The related fields of momentum and reversal in financial asset returns have proved to be an important area in the literature. However, existing work does not evaluate the profitability of intraday reversal and relative strength. This is a large gap in our understanding, given that the time horizon over which portfolios of winners and losers are formed proved vital in establishing the two effects. Chapter 4 provides a detailed study in this area, using high frequency data. The sample is

very large, encompassing all of the S&P 500 constituents and spanning the whole of 2005. Looking at a range of different formation and holding periods for portfolios over 10-60 minutes, it was found that there is a clear intraday reversal effect in returns. This finding was shown to be robust across both months of the year and days of the week. Results accord with our knowledge of intraday trading activity, as the reversal effect is most pronounced at the start and end of the trading day. Further, it is conditional on size, with the largest and most actively traded stocks showing the greatest returns from buying losers and selling winners. These results are important for traders, and in particular day traders. Survey evidence suggests that this group primarily act as momentum traders; this is not a profitable strategy and is likely to be a factor in why day traders often exhaust their capital, and exit the market, in a short period of time.

Finally, Chapter 5 examines one of the longest standing forms of technical analysis, point and figure charting. The point and figure technique has been used for over 100 years. However, there has been scant academic investigation in this area. The empirical work in this thesis investigated the profitability of a point and figure trading strategy using an extremely large sample, constituting the entire contents of the NYSE consolidated trades database for 2005. This 'real-time' data was the original intended application of point and figure. The results show that the charting method forms a useful filtering tool, but also that profits are available to traders. Clear trading rules are well-established for point and figure charting, and these were investigated. The least restrictive rules could not be profitably employed by traders; however, the more complex rules did generate significant returns. These results run counter to weak-form market efficiency.

These four distinct empirical chapters represent a large extension of our knowledge of technical trading strategies. Significant limitations and gaps in the existing literature have been addressed and new, innovative developments have been made. Whilst the findings have important implications in terms of efficient markets, they

have wider significance. Specifically, to the large number of market practitioners who actively employ technical analysis on a day-to-day basis.

Bibliography

- S. B. Achelis. Technical Analysis from A to Z. McGraw-Hill, New York, 2001.
- S. S. Alexander. Price Movements in Speculative Markets: Trends or Random Walks. Industrial Management Review, II:7–26, 1961.
- S. S. Alexander. Price movements in speculative markets: Trends of random walks 2. In P. Cootner, editor, The Random Character of Stock Market Prices. MIT Press, Cambridge, Mass., 1964.
- J. A. Anderson and R. W. Faff. Point and figure charting: A computational methodology and trading rule performance in the S&P 500 futures market. International Review of Financial Analysis, 17(1):198–217, 2008.
- A. Antoniou, E. Galariotis, S. Spyrou, and M. Lane. Short-term contrarian strategies in the London Stock Exchange: are they profitable? which factors affect them? Journal of Business Finance & Accounting, 33(5-6):839–867, 2006.
- T. Arnsward. Investment behaviour of german equity fund managers: An exploratory analysis of survey data. Discussion paper 08/01, Economic Research Centre of the Deutsche Bundesbank, 2001.
- B. Barber and T. Odean. The courage of misguided convictions. Financial Analysts Journal, 55(6):41–55, 1999.
- N. Barberis, A. Shleifer, and R. Vishny. A model of investor sentiment. Journal of Financial Economics, 49(3):307–343, 1998.
- S. Berkowitz, D. Logue, and E. Noser. The total cost of transactions on the NYSE. The Journal of Finance, 43(1):97–112, 1988.
- H. Bessembinder and K. Chan. Market Efficiency and the Returns to Technical Analysis. Financial Management, 27(2):5–17, 1998.
- P. J. W. N. Bird. The weak form efficiency of the London Metal Exchange. Applied Economics, 17(5):571–588, 1985.
- W. Brock, J. Lakonishok, and B. LeBaron. Simple Technical Trading Rules and the Stochastic Properties of Stock Returns. Journal of Finance, 47(5):1731–1764, 1992.

- M. Brockmann, T. Gasser, and E. Herrmann. Locally adaptive bandwidth choice for kernel regression estimators. Journal of the American Statistical Association, 88(24):1302–1309, 1993.
- B. Brorsen and S. Irwin. Futures funds and price volatility. The Review of Futures Markets, 6:118–135, 1987.
- C. Brownlees and G. Gallo. Financial econometric analysis at ultra-high frequency: Data handling concerns. Computer Statistics and Data Analysis, 51:2232–2245, 2006.
- G. Bry and C. Boschan. Cyclical analysis of time series: Selected procedures and computer programs. Technical Paper 20, National Bureau of Economic Research (Distributed by Columbia University Press), 1971.
- T. N. Bulkowski. Encyclopedia of chart patterns. Wiley trading series. John Wiley and Sons, Hoboken, New Jersey, 2 edition, 2005.
- M. M. Carhart. On persistence in mutual fund performance. The Journal of Finance, 52(1):57–82, 1997.
- B. Chakrabarty and C. Trzcinka. Momentum: Does the database make a difference? Journal of Financial Research, 29(4):441–462, 2006.
- K. Chan, A. Hameed, and W. Tong. Profitability of momentum strategies in the international equity markets. The Journal of Financial and Quantitative Analysis, 35(2):153–172, 2000.
- L. K. C. Chan, N. Jegadeesh, and J. Lakonishok. Momentum strategies. The Journal of Finance, 51(5):1681–1713, 1996.
- L. K. C. Chan, N. Jegadeesh, and J. Lakonishok. The profitability of momentum strategies. Financial Analysts Journal, 55(6):80–90, 1999.
- P. H. K. Chang and C. L. Osler. Head and Shoulders: Not just a Flaky Pattern. Federal Reserve Bank of New York: Staff Reports, (4), 1995.
- P. H. K. Chang and C. L. Osler. Methodical Madness: Technical Analysis and the Irrationality of Exchange-Rate Forecasts. Economic Journal, 109(458):636–661, 1999.
- Y.-W. Cheung and M. Chinn. Currency traders and exchange rate dynamics: A survey of the us market. Journal of International Money and Finance, 20:439–471, 2001.
- Y.-W. Cheung and C. Y.-P. Wong. A survey of market practitioners' views on exchange rate dynamics. Journal of International Economics, 51(2):401–419, 2000.
- Y.-W. Cheung, M. Chinn, and I. Marsh. How do uk-based foreign exchange dealers think their market operates? NBER Working Paper, (7524), 2000.

- Y.-W. Cheung, M. D. Chinn, and I. W. Marsh. How do uk-based foreign exchange dealers think their market operates? International Journal of Finance & Economics, 9(4):289–306, 2004.
- C. K. Chu and J. S. Marron. Choosing a kernel regression estimator. Statistical Science, 6(4):404–419, 1991.
- P. Clay. Forecasting security prices. volume 20, pages 244–249. American Statistical Association, 1925.
- J. Conrad and G. Kaul. An anatomy of trading strategies. Rev. Financ. Stud., 11(3): 489–519, 1998.
- C. J. Corrado and S.-H. Lee. Filter rule tests of the economic significance of serial dependencies in daily stock returns. Journal of Financial Research, 15(4):369–390, 1992.
- A. Cowles. Can stock market forecasters forecast? Econometrica: Journal of the Econometric Society, 1(3):309–324, 1933.
- S. W. Cunningham. The Predictability of British Stock Market Prices. Applied Statistics, 22(3):315–31, 1973.
- K. Daniel, D. Hirshleifer, and A. Subrahmanyam. Investor psychology and security market under- and overreactions. Journal of Finance, 53(6):1839–1885, 1998.
- G. D’Avolio. The market for borrowing stock. Journal of Financial Economics, 66: 271–306, 2002.
- E. R. Dawson and J. M. Steeley. On the existence of visual technical patterns in the UK stock market. Journal of Business Finance and Accounting, 30(1-2):263–293, 2003.
- W. F. M. De Bondt and R. Thaler. Does the stock market overreact? The Journal of Finance, 40(3):793–805, 1985.
- W. F. M. De Bondt and R. Thaler. Further evidence of investor overreaction and stock market seasonality. Journal of Finance, 42:557–581, 1987.
- V. DeVilliers and O. Taylor. The Point & Figure method of anticipating stock price movements. Windsor Books, New York, 1933.
- F. X. Diebold and J. A. Nason. Nonparametric exchange rate prediction. Journal of International Economics, 28:315–332, 1990.
- T. J. Dorsey. Point & Figure Charting: The Essential Application for Forecasting and Tracking Market Prices. Wiley, New Jersey, 2007.
- T. J. Dorsey, T. F. DeRosier, P. L. Keeton, S. L. Morrison, and J. B. Parker. Commodity Strategies: High-Profit Techniques for Investors and Traders. Wiley, New Jersey, 2007.

- M. Dryden. Filter Tests of UK Share Prices. Applied Economics Letters, 1(4): 261–75, 1970.
- J. Du Plessis. The Definitive Guide to Point and Figure. Harriman House Publishing, Petersfield, 2005.
- R. Edwards and J. Magee. Technical Analysis of Stock Trends. American Management Association, 8 edition, 2001.
- R. D. Edwards and J. Magee. Technical Analysis of Stock Trends. Stock Trend Service, Springfield, Mass., 1 edition, 1948.
- R. J. Elliott and J. Hinz. Portfolio optimization, hidden markov models, and technical analysis of p&f-charts. International Journal of Theoretical and Applied Finance, 5(4):385–399, 2002.
- C. A. Ellis and S. A. Parbery. Is smarter better? a comparison of adaptive and simple moving average trading strategies. Research in International Business and Finance, 19:399–411, 2005.
- R. L. Eubank. Nonparametric Regression and Spline Smoothing. Marcel Dekker, New York, second edition, 1999.
- F. J. Fabozzi, C. K. Ma, W. T. Chittenden, and R. D. Pace. Predicting intraday price reversals. Journal of Portfolio Management, 21(2):42–54, 1995.
- E. Fama and K. French. Value versus growth: The international evidence. Journal of Finance, 53:1975–1999, 1998.
- E. F. Fama and M. E. Blume. Filter Rules and Stock-Market Trading. Journal of Business, 39(1, Part 2: Supplement on Security Prices):226–241, 1966.
- I. Figelman. Stock Return Momentum and Reversal. The Journal of Portfolio Management, 34(1):51–67, 2007.
- W. M. Fong and L. H. M. Yong. Chasing trends: recursive moving average trading rules and internet stocks. Journal of empirical finance, 12:43–76, 2005.
- F. Foster and S. Viswanathan. A theory of intraday variation in volumes, variance, and trading costs in securities markets. Review of Financial Studies, 3:593–624, 1990.
- J. Frankel and K. Froot. Chartists, fundamentalists, and trading in the foreign exchange... American Economic Review, 80(2):181–185, 1990.
- K. French. Stock Returns and the Weekend Effect. Journal of Financial Economics, 8(1):55–69, 1980.
- A. K.-W. Fung, D. M. Y. Mok, and K. Lam. Intraday price reversals for index futures in the us and hong kong. Journal of Banking & Finance, 24(7):1179–1201, 2000.

- W. Gann. New Stock Trend Detector: A Review of the 1929-1932 Panic and the 1932-1935 Bull Market with Rules and Charts for Detecting Trend of Stocks. Traders Press, 1936.
- T. Gasser and H. G. Muller. Smoothing techniques for curve estimation. In Lecture Notes in Mathematics, chapter Kernel estimation of regression functions, pages 23–68. Springer, New York, 1979.
- T. Gasser, A. Kneip, and W. Kohler. A flexible and fast method for automatic smoothing. Journal of the American Statistical Association, 86(415):643–652, 1991.
- T. Gehrig and L. Menkhoff. Extended evidence on the use of technical analysis in foreign exchange. International Journal of Finance & Economics, 11(4):327–338, 2006.
- R. Gencay. Linear, non-linear and essential foreign exchange rate prediction with simple technical trading rules. Journal of International Economics, 47:91–107, 1999.
- R. L. Giles. Forward looking uk interest rates: a point and figure chart approach. Managerial Finance, 31(5):1–13, 2005.
- J. L. Grant, A. Wolf, and S. Yu. Intraday price reversals in the us stock index futures market: A 15-year study. Journal of Banking and Finance, 29, 2005.
- J. Griffin, J. Harris, and S. Topaloglu. The dynamics of institutional and individual trading. Journal of Finance, pages 2285–2320, 2003a.
- J. M. Griffin, Xiuqing Ji, and J. S. Martin. Momentum investing and business cycle risk: Evidence from pole to pole. Journal of Finance, LVIII(6):2515–2547, 2003b.
- Group of Thirty. The Foreign Exchange Market in the 1980s. Group of Thirty, New York, 1985.
- R. Gutierrez and E. Kelley. The Long-Lasting Momentum in Weekly Returns (Digest Summary). Journal of Finance, 63(1):415–447, 2008.
- A. Hameed and Y. Kusnadi. Momentum strategies: Evidence from pacific basin stock markets. The Journal of Financial Research, XXV(3):383–397, 2002.
- W. P. Hamilton. The Stock Market Barometer. Harper and Row Brothers, New York, 1922. republished in 1998 by John Wiley and Sons.
- W. Hardle. Applied Nonparametric Regression. Cambridge University Press, Cambridge, 1990.
- L. Harris. A transaction data study of weekly and intradaily patterns in stock returns. Journal of Financial Economics, 161(1):99–117, 1986.
- K. Hauschild and M. Winkelmann. Kapitalmarkteffizienz und Point & Figure Analyse. Kredit und Kapital, 18:240–264, 1985.

- E. Herrmann. Local bandwidth choice in kernel regression estimation. Journal of Computational and Graphical Statistics, 6(1):35–54, 1997.
- E. Hol and J. Koopman. Stock index volatility forecasting with high frequency data. Tinbergen Institute Discussion Papers, 02-068/4, 2002.
- H. Hong and J. C. Stein. A unified theory of underreaction, momentum trading, and overreaction in asset markets. Journal of Finance, 54(6):2143–2184, 1999.
- H. Hong, T. Lim, and J. C. Stein. Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies. Journal of Finance, 55:265–295, 2000.
- R. Hudson, M. Dempsey, and K. Keasey. A Note on the Weak Form Efficiency of Capital Markets: The Application of Simple Technical Trading Rule to the UK Stock Prices – 1935 to 1994. Journal of Banking and Finance, 20(6):1121–32, 1996.
- S. Hvidkjaer. A Trade-Based Analysis of Momentum. Review of Financial Studies, 19(2):457–491, 2006.
- F. E. J. James. Monthly Moving Averages—An Effective Investment Tool? Journal of Financial and Quantitative Analysis, 3(3, Special Issue: Random Walk Hypothesis):315–326, 1968.
- N. Jegadeesh. Evidence of predictable behavior of security returns. Journal of Finance, 45(3):p881 – 898, 1990.
- N. Jegadeesh and S. Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. Journal of Finance, 48:65–91, 1993.
- N. Jegadeesh and S. Titman. Overreaction, Delayed Reaction, and Contrarian Profits. Review of Financial Studies, 8(4):973–93, 1995. Downloaded: Y.
- N. Jegadeesh and S. Titman. Profitability of momentum strategies: An evaluation of alternative explanations. Journal of Finance, 56:699–720, 2001.
- N. Jegadeesh, J. Kim, S. D. Krische, and C. M. C. Lee. Analyzing the analysts: When do recommendations add value? The Journal of Finance, 59(3):1083–1124, 2004.
- M. C. Jensen and G. A. Benington. Random Walks and Technical Theories: Some Additional Evidence. Journal of Finance, 25(2, Papers and Proceedings of the Twenty-Eighth Annual Meeting of the American Finance Association New York, N.Y. December, 28-30, 1969):469–482, 1970.
- C. Jones. A century of stock market liquidity and trading costs. Working Paper, Graduate School of Business, Columbia University, 2002.
- M. C. Jones, S. J. Davies, and B. U. Park. Versions of kernel-type regression estimators. Journal of the American Statistical Association, 89(427):825–832, 1994.
- D. J. Jordan and J. D. Diltz. The profitability of day traders. Financial Analysts Journal, 59(6):85–94, 2003.

- P. J. Kaufman. New trading systems and methods. John Wiley & Sons, 4 edition, 2005.
- K. A. Kavajecz and E. R. Odders-White. Technical analysis and liquidity provision. Review of Financial Studies, 17(4):1043–1071, 2004.
- M. G. Kendall. The analysis of economic time series. i. Journal of the Royal Statistical Society, A(CXVI):1125, 1953.
- W. I. King. Technical methods of forecasting stock prices. Journal of the American Statistical Association, 29(187):323–325, 1934.
- R. La Porta, J. Lakonishok, A. Shleifer, and R. Vishny. Good news for value stocks: Further evidence on market efficiency. Journal of Finance, 52:859–874, 1997.
- J. Lakonishok and E. Maberly. The Weekend Effect: Trading Patterns of Individual and Institutional Investors. Journal of Finance, 45(1):231–243, 1990.
- J. Lakonishok, A. Shleifer, and R. Vishny. Contrarian investment, extrapolation, and risk. Journal of Finance, 49:1541–78, 1994.
- V. W. Lam, T. T. Chong, and W. K. Wong. Profitability of intraday and interday momentum strategies. Applied Economics Letters, 14(15):1103–1108, 2007.
- C. Lee and I. Mathur. Trading rule profits in European currency spot cross-rates. Journal of Banking and Finance, 20(5):949–962, 1996a.
- C. Lee and I. Mathur. A comprehensive look at the efficacy of technical trading rules applied to cross-rates. The European Journal of Finance, 2(4):389–411, 1996b.
- C. I. Lee, K. C. Gleason, and I. Mathur. Trading Rule Profits in Latin American Currency Spot Rates. International Review of Financial Analysis, 10(2):135–56, 2001.
- C. M. C. Lee and B. Swaminathan. Price momentum and trading volume. Journal of Finance, 55(5):2017–2069, 2000.
- B. N. Lehmann. Fads, Martingales, and Market Efficiency. The Quarterly Journal of Economics, 105(1):1–28, 1990.
- R. M. Levich and L. R. I. Thomas. The significance of technical trading-rule profits in the foreign exchange market: a bootstrap approach. Journal of International Money and Finance, 12(5):451–474, 1993.
- R. A. Levy. Relative Strength as a Criterion for Investment Selection. Journal of Finance, 22(4):595–610, 1967.
- J. Liew and M. Vassalou. Can book-to-market, size and momentum be risk factors that predict economic growth? Journal of Financial Economics, 57(2):221–245, 2000.

- A. W. Lo and A. C. MacKinlay. When Are Contrarian Profits Due to Stock Market Overreaction? Review of Financial Studies, 3(2):175–205, 1990.
- A. W. Lo, H. Mamaysky, and J. Wang. Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation. Journal of Finance, 55(4):1705–65, 2000.
- D. E. Logue, R. J. Sweeney, and R. D. Willett. The Speculative Behavior of Foreign Exchange Rates During the Current Float. Journal of Business Research, 6:159–74, 1978.
- B. Lucke. Are technical trading rules profitable? evidence for head-and-shoulder rules. Applied Economics, 35(1):33–40, 2003.
- Y.-H. Lui and D. Mole. The use of fundamental and technical analyses by foreign exchange dealers: Hong kong evidence. Journal of International Money and Finance, 17(3):535–545, 1998.
- Mackinson Cowell. Stock lending - a perspective. Report, 2005.
- B. Maillet and T. Michel. Further insights on the puzzle of technical analysis profitability. The European Journal of Finance, 6(2):196–224, 2000.
- B. G. Malkiel. A random walk down Wall Street : including a life-cycle guide to personal investing. W.W. Norton, New York ; London, rev. and updated edition, 1999.
- B. Marshall, M. Young, and R. Cahan. Are candlestick technical trading strategies profitable in the Japanese equity market? Review of Quantitative Finance and Accounting, 31:191–207, 2008a.
- B. R. Marshall, M. R. Young, and L. C. Rose. Candlestick technical trading strategies: Can they create value for investors? Journal of Banking and Finance, 30(8):2303–2323, 2006.
- B. R. Marshall, R. H. Cahan, and J. M. Cahan. Does intraday technical analysis in the u.s. equity market have value? Journal of Empirical Finance, 15:199–210, 2008b.
- A. Martin. Technical trading rules in the spot foreign exchange markets of developing countries. Journal of Multinational Financial Management, 11(1):59–68, 2001.
- J. S. Martin and B. G. Grundy. Understanding the nature of the risks and the source of the rewards to momentum investing. Review of Financial Studies, 14(1):29–78(50), 2001.
- A. K. Meese, Richard A. Rose. Nonlinear, nonparametric, nonessential exchange rate estimation. American Economic Review, 80(2):192–196, 1990.
- L. Menkhoff. Examining the use of technical currency analysis. International Journal of Finance & Economics, 2(4):307–318, 1997.

- L. Menkhoff and U. Schmidt. the use of trading strategies by fund managers: some first survey evidence. Applied Economics, 37:1719–1730, 2005.
- T. C. Mills. Technical Analysis and the London Stock Exchange: Testing Trading Rules Using the FT30. International Journal of Financial Economics, 2(7?):319–331, 1997.
- B. Mizrach and S. Weerts. Experts online: An analysis of trading activity in a public Internet chat room. Journal of Economic Behavior and Organization, 2009.
- T. J. Moskowitz and M. Grinblatt. Do industries explain momentum? Journal of Finance, 54(4):1249–1290, 1999.
- J. J. Murphy. Technical Analysis of the Financial Markets. New York Institute of Finance, Paramus, NJ, 1999.
- E. A. Nadaraya. On estimating regression. Theory of Probability and its Applications, 10:186–190, 1964.
- S. Nelson. The ABC of Stock Market Speculation. Fraser Publishing, New York, 1903.
- S. Nison. Beyond candlesticks: New Japanese charting techniques revealed. Wiley, Hoboken, New Jersey, 1994.
- S. Nison. Japanese Candlestick Charting. New York Institute of Finance (Prentice Hall Press), Paramus, NJ, 2001.
- J. Nofsinger and R. Sias. Herding and feedback trading by institutional and individual investors. Journal of Finance, 54(6):2263–2295, 1999.
- T. Oberlechner. Importance of technical and fundamental analysis in the European foreign exchange market. International Journal of Finance & Economics, 6(1): 81–93, 2001.
- T. Odean. Volume, volatility, price, and profit when all traders are above average. Journal of Finance, 53(6):1887–1934, 1998.
- D. Olson. Have trading rule profits in the currency markets declined over time? Journal of Banking and Finance, 28(1):85–105, 2004.
- B. U. Park and J. S. Marron. Comparison of data-driven bandwidth selectors. Journal of the American Statistical Association, 85(409):66–72, 1990.
- C.-H. Park and S. H. Irwin. What do we know about the profitability of technical analysis? Journal of Economic Surveys, 21(4):786–826, 2007.
- T. Poggio and D. Beymer. Early Visual Learning, chapter Regularization Networks for Visual Learning. Oxford University Press, Oxford, 1996.

- M. Pring. Technical Analysis Explained: The Successful Investor's Guide to Spotting Investment Trends and Turning Points. McGraw-Hill, New York, 3 edition, 1985.
- M. Pring. Martin Pring's introduction to technical analysis. McGraw-Hill, New York, 1998.
- M. Ratner and R. P. C. Leal. Tests of technical trading strategies in the emerging equity markets of Latin America and Asia. Journal of Banking and Finance, 23 (12):1887–1905, 1999.
- M. J. Ready. Profits from Technical Trading Rules. Financial Management, 31(3): 43–61, 2002.
- M. Reinganum. The anomalous stock market behavior of small firms in january – empirical tests for tax-loss selling effects. Journal of Financial Economics, 12: 89–104, 1983.
- R. Rhea. The Dow Theory. Barron's, New York, 1932. republished in 1993 by Fraser Publishing Company.
- J. Ritter. The buying and selling behavior of individual investors at the turn of the year. Journal of Finance, 43(3):701–717, 1988.
- H. V. Roberts. Stock-market “patterns” and financial analysis: Methodological suggestions. Journal of Finance, 14(1):1–10, 1959.
- K. G. Rouwenhorst. International momentum strategies. Journal of Finance, 53: 267–284, 1998.
- K. G. Rouwenhorst. Local return factors and turnover in emerging stock markets. Journal of Finance, 54(4):1439–1464, 1999.
- M. Rozeff and W. Kinney. Capital market seasonality: The case of stock returns. Journal of Financial Economics, 3:379–402, 1976.
- P. A. Samuelson. Proof that properly anticipated prices fluctuate randomly. Industrial Management Review, 6:41–49, 1965.
- D. Sauer and C. Chen. Mean Reversion in the United Kingdom Stock Market and its Implications for a Profitable Trading Strategy. Journal of Business Finance and Accounting, 23(9-10):1379–95, 1996.
- G. Savin, P. Weller, and J. Zvingelis. The predictive power of head-and-shoulders price patterns in the U.S stock market. Journal of Financial Econometrics, 5(2): 243–265, 2007.
- R. Schabacker. Technical Analysis and Stock Market Profits. Harriman House, Petersfield, Hampshire, 2005.
- R. Schabacker. Technical analysis and market profits. Sarchmont, N.Y., c1932.

- E. G. Schaefer. Technical studies correlated with Dow theory long-term and intermediate trends. Dow Theory Trader, Washington, 1956.
- S. Schulmeister. Profitability of technical stock trading: Has it moved from daily to intraday data? Review of Financial Economics, 2008.
- S. Sheather and M. C. Jones. A reliable data-based bandwidth selection method for kernel density estimation. Journal of the Royal Statistical Society, 53(3):683–690, 1991.
- W. Silber. Technical trading: When it works and when it doesn't. The Journal of Derivatives, 1(3):39–44, 1994.
- B. W. Silverman. Density Estimation for Statistics and Data Analysis. Chapman and Hall, London, 1986.
- S. Smidt. Amateur Speculators. Graduate School of Business and Public Administration, Cornell University, Ithaca, NY, 1965.
- L. Stevens. Essential Technical Analysis. John Wiley and Sons, New York, 2002.
- B. Stewart. An analysis of speculative trading in grain futures. Technical Bulletin 1001, US Department of Agriculture, Washington, D.C., 1949.
- R. Stottner. P&F-Filteranalyse, Averaging-Strategie und Buy&Hold-Anlageregel. Jahrbücher für Nationalökonomie und Statistik, 207(4):374–390, 1990.
- R. Sullivan, A. Timmerman, and H. White. Data-snooping, Technical Trading Rule Performance, and the Bootstrap. Journal of Finance, 54(5):1647–91, 1999.
- R. J. Sweeney. Beating the Foreign Exchange Market. Journal of Finance, 41(1): 163–182, 1986.
- R. J. Sweeny and P. Surajaras. The stability of speculative profits in the foreign exchanges. In R. M. C. Guimaraes, B. G. Kingsman, and S. J. Taylor, editors, A Reappraisal of the Efficiency of Financial Markets. Springer-Verlag, Heidelberg, 1989.
- M. P. Taylor and H. Allen. The use of technical analysis in the foreign exchange market. Journal of International Money and Finance, 11(3):304–314, 1992.
- S. Thomas. Discussion of short sales constraints and momentum in stock returns. Journal of Business Finance and Accounting, 33(3):616–631, 2006.
- Thomson-Financial. Datastream International research database, 2005.
- G. S. Watson. Smooth regression analysis. Sankhya, 26:359–372, 1964.
- H. Weber and K. Zieg. Complete Guide to Point-and-Figure Charting. Harriman House Publishing, 2003.

- R. E. Weintraub. On speculative prices and random walks a denial. Journal of Finance, 18(1):59–66, 1963.
- R. Wood, T. McInish, and J. Ord. An investigation of transactions data for nyse stocks. Journal of Finance, 40:723–9, 1985.
- S. Yu, J. Rentzler, and A. Wolf. Nasdaq-100 index futures: Intraday momentum or reversal? Journal of Investment Management, 3(3), 2005.
- K. C. Zieg and P. J. Kaufman. Point-and-Figure commodity trading techniques. Investors Intelligence, Larchmont, NY, 1975.