

**The Technical Analysis Method of Moving Average Trading:
Rules That Reduce the Number of Losing Trades**

Marcus C. Toms

A thesis submitted for the degree of Doctor of Philosophy
Department of Electrical, Electronic and Computer Engineering

Newcastle University

August 2011

Abstract

A general issue with moving average trading is the assumption that all buy/sell signals result in a trading action. The argument that such trading rules are representative of trading practice is highly questionable. This thesis proposes two new moving average trading rules designed to capture trading practice. The first trading rule is the trade reduction rule and is based on the idea of allowing a trade to run. The second trading rule is the positive autocorrelation rule and is based on the idea of only trading if it is believed to be profitable to do so. The trading rules are tied to moving average trading via the buy/sell signal generating mechanism and alter the way the price crossover rule responds to the buy/sell signals. Simulations of portfolios of UK equities find that the trading rules uncover information that is missed by the price crossover rule and there is evidence that this information is financially exploitable. This motivates the argument that the information needed for trading to be economically viable is observable in the price. The trading rules also establish a link with the market microstructure literature. The trading rules uncover issues of informed trading (asymmetric information), liquidity, adverse selection and price impact. The strongest interpretation that can be applied to the trading rules in this context is that they are examples of informed trading. Compared to the price crossover rule, the trading rules are better able to extract meaning from or are better able to understand the same price information.

Acknowledgements

I would like to thank my supervisor Dr Graeme Chester for his continuing support. I would also like to thank my advisor Professor Robert Hudson for offering words of encouragement at the right time. Similar thanks go to my family. I would particularly like to thank my father whose help in getting through the final year was invaluable. Special thanks are for Yu-Chun and Marc for memories that will last a lifetime.

Table of Contents

	List of Figures	iv
	List of Tables	viii
1	Introduction	1
1.1	Background	1
1.2	Moving average trading	2
1.3	Efficient market hypothesis	7
1.4	Moving average trading and market efficiency	10
1.5	Motivation	12
1.6	Remodelling the price crossover rule	12
1.7	Research objectives	17
1.8	Scope	18
1.9	Contributions	18
1.10	Thesis outline	19
2	Literature Review	20
2.1	Technical analysis	20
2.2	Survey literature	23
2.3	Conclusion	27
3	Data and Test Method	29
3.1	Data	29
3.1.1	Requirements and considerations	30
3.1.2	Content and format	31
3.1.3	Pre-processing the close	32

3.1.4	Reconstructing the spread	33
3.1.5	Portfolios and sub-periods	41
3.2	Simulation issues	45
3.3	Test method	46
3.4	Summary	49
4	Trade Reduction Rule	51
4.1	Trade reduction rule	52
4.1.1	Bounded moving average	52
4.1.2	Trade reduction algorithm	55
4.2	Results for the mean return per trade	68
4.2.1	Long	69
4.2.2	Short	79
4.3	Conclusion	89
5	Positive Autocorrelation Rule	93
5.1	Positive autocorrelation rule	94
5.1.1	Positive autocorrelation and profitability	95
5.1.2	Variance ratio tests	98
5.1.3	Variance ratio profiles	103
5.1.4	Positive autocorrelation rule	105
5.2	Results	111
5.2.1	Data	112
5.2.2	Simulation method	113
5.2.3	Illustrative example	115
5.2.4	Portfolio simulations	119
5.3	Conclusion	124

6	Conclusion and Further Work	127
6.1	Further work	132
	Appendix	137
A.1	FTSE 100 Portfolio (Chapters 3 and 4)	137
A.2	FTSE 250 Portfolio (Chapters 3 and 4)	140
A.3	FTSE Small Cap Portfolio (Chapters 3 and 4)	145
A.4	FTSE Fledgling Portfolio (Chapters 3 and 4)	149
A.5	Sharescope	152
A.6	Updata Technical Analyst	153
A.7	FTSE 100 Portfolio (Chapter 5)	153
A.8	FTSE 250 Portfolio (Chapter 5)	154
A.9	FTSE Small Cap Portfolio (Chapter 5)	155
A.10	FTSE Fledgling Portfolio (Chapter 5)	156
A.11	$VP(25,1) = (\mathbf{VP}(25) - \overline{\mathbf{RW}}) \mathbf{e}_1$	157
	References	158

List of Figures

1.1	Example of the moving average trading process	3
3.1	Mean annual bid-ask spreads	39
4.1	Example of the moving average and its mapping to the bounded moving average	54
4.2	Example frequency distribution of the position of minima and maxima	55
4.3	Example trades for the long only trade reduction rule in terms of the path of the bounded moving average	61
4.4a	Number of winning and losing trades for the long only price crossover and trade reduction rules for the FTSE 100 portfolio for the period 01-January-1965 to 30-June-2009	65
4.4b	Number of winning and losing trades for the short only price crossover and trade reduction rules for the FTSE 100 portfolio for the period 01-January-1965 to 30-June-2009	66
4.5	Sum of the winning and losing trades for the long only and short only price crossover and trade reduction rules for FTSE 100 portfolio for the period 01-January-1965 to 30-June-2009	67
4.6a	Mean return per trade for the long only price crossover and trade reduction rules for the FTSE 100 portfolio	71
4.6b	Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the long only price crossover and trade reduction rules for the FTSE 100 portfolio	72
4.7a	Mean return per trade for the long only price crossover and trade reduction	

	rules for the FTSE 250 portfolio	73
4.7b	Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the long only price crossover and trade reduction rules for the FTSE 250 portfolio	74
4.8a	Mean return per trade for the long only price crossover and trade reduction rules for the FTSE Small Cap portfolio	75
4.8b	Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the long only price crossover and trade reduction rules for the FTSE Small Cap portfolio	76
4.9a	Mean return per trade for the long only price crossover and trade reduction rules for the FTSE Fledgling portfolio	77
4.9b	Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the long only price crossover and trade reduction rules for the FTSE Fledgling portfolio	78
4.10a	Mean return per trade for the short only price crossover and trade reduction rules for the FTSE 100 portfolio	81
4.10b	Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the short only price crossover and trade reduction rules for the FTSE 100 portfolio	82
4.11a	Mean return per trade for the short only price crossover and trade reduction rules for the FTSE 250 portfolio	83
4.11b	Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the short only price crossover and trade reduction rules for the FTSE 250 portfolio	84
4.12a	Mean return per trade for the short only price crossover and trade reduction	

	rules for the FTSE Small Cap portfolio	85
4.12b	Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the short only price crossover and trade reduction rules for the FTSE Small Cap portfolio	86
4.13a	Mean return per trade for the short only price crossover and trade reduction rules for the FTSE Fledgling portfolio	87
4.13b	Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the short only price crossover and trade reduction rules for the FTSE Fledgling portfolio	88
5.1	Artificial price series that exhibit positive autocorrelation, negative autocorrelation and random walk dynamics and the return to trading the price series using a 10-day price crossover rule	97
5.2	Artificial price series that exhibit positive autocorrelation, negative autocorrelation and random walk dynamics and their variance ratios	101
5.3	The first three eigenvectors for $\mathbf{VP}(100)$ estimated using 1000 random walk time series	105
5.4	The first three eigenvectors for $\mathbf{VP}(25)$ estimated using 1000 random walk time series	107
5.5	Variance explained by the first 10 principal components for $\mathbf{VP}(25)$	108
5.6	Cumulative frequency distribution for $VP(25,1)$ estimated using another 1000 random walk time series	109
5.7	Equity curves for the 10-day moving average price crossover and positive autocorrelation rules for Yule Catto	117
5.8	Return for the price crossover and positive autocorrelation rules for Yule Catto	118
5.9	Return for the price crossover and positive autocorrelation rules for the FTSE	

	100, FTSE 250, FTSE Small Cap and FTSE Fledgling portfolios	120
5.10	Ratio of the capital needed to fund trading for the price crossover and positive autocorrelation rules	124
6.1	Trading and the trader's perception of the relationship in information flow	129

List of Tables

1.1	Mean return per trade for Yule Catto	14
3.1	Mean bid-ask spreads	38
3.2	Number of stocks in each portfolio and sub-period	44
5.1	Simulation parameters	115

Chapter 1

Introduction

A *trading system* is a systematic method for buying and selling financial instruments with a view to consistently making money. As such, not only do trading systems require that prices are predictable but also that the predictable component is financially exploitable. This thesis considers the technical analysis method of *moving average trading* as the basis for a simple stock trading system. The question asked is whether remodelling the trading rules to reduce the number of losing trades increases the mean return per trade to the extent that the trading rules are profitable and, if so, whether this is economically significant.

1.1 Background

Technical analysis is an approach to predicting future price movements based on identifying patterns in prices, volume and other market statistics. Technical analysis usually proceeds by recording market activity in graphical form and then deducing the probable future trend from the pictured history. The premise is that prices exhibit various geometric regularities, which, once identified, inform the trader what is likely to happen next. This in turn allows the trader to run a profitable trading strategy. Technical analysis is prevalent in financial markets and is readily accessible in practitioner texts such as Pring (2002), in the form of tools provided by online brokers such as Barclays Stockbrokers (www.stockbrokers.barclays.co.uk) as well as in the form of commentary in the financial and investment press.

Of the technical analysis methods studied in the literature, it is moving average trading that is perhaps the most compelling. Moving average trading has been shown to uncover predictable behaviour in the first and second moments of the returns distribution and this result has since been replicated for many different markets and asset classes. However, there is little evidence that this implies a market beating trading strategy. Taylor (2005) discusses moving average trading in detail. The conclusion is that while there is plenty of evidence that moving average trading has been able to uncover predictable behaviour in the returns distribution and where this has sometimes been sufficiently precise to allow risk adjusted profits of several per cent per annum, the successful application of moving average trading is, in general, restricted to prices before the 1990's.

An issue that arises in response to this is that the moving average trading rules discussed in Taylor (2005) have not changed since their introduction in Brock *et al.* (1992). The trading rules suffer from various problems, the most significant of which is that they are either not profitable or borderline profitable. By profitable, it is meant that the mean return per trade is greater than the expenditure on costs such that, on average, each trade earns the mean return per trade minus costs. The background to this thesis is to ask whether remodelling the trading rules to reduce the number of losing trades increases the mean return per trade to the extent that the trading rules are profitable and, if so, whether this is economically significant.

1.2 Moving average trading

Moving average trading refers to the practice of systematically buying and selling whenever the price crosses its average. The idea is that prices move in trends such that at each point in time the price is either in an *uptrend* or in a *downtrend*. An uptrend is defined as a period of

rising prices and a downtrend is defined as a period of falling prices. When the price cuts up through its average from below, because recent prices are higher than older prices, the price is said to be in an uptrend and a *buy signal* occurs. Similarly, when the price cuts down through its average from above, because recent prices are lower than older prices, the price is said to be in a downtrend and a *sell signal* occurs. The response following a buy signal is to buy and the response following a sell signal is to sell. If the change in the price level in between buy and sell signals is sufficient to cover costs, moving average trading is profitable. Conversely, if the change in the price level in between buy and sell signals is not sufficient to cover costs, moving average trading is loss making. Figure 1.1 plots an example.




Figure 1.1 Example of the moving average trading process.

Moving average trading is usually defined in the form of a *trading rule*. A trading rule is a numerical method that maps the price onto investment decisions. A typical decision variable at time t is the quantity of an asset q_{t+1} that is held from the time of the price observation at

time t until the time of the next price observation at time $t + 1$. Let $ma_t^n(p_t)$ denote the n -day moving average of the price p_t at time t :

$$ma_t^n(p_t) = \frac{1}{n} \sum_{i=0}^{n-1} p_{t-i} \quad (1.1)$$

The simplest moving average trading rule is the *price crossover rule*:

$$q_1 = 0$$

///// RULE

$$q_{t+1} = q_t$$

$$\text{IF } (p_t > ma_t^n(p_t)) \text{ THEN } q_{t+1} = 1$$

$$\text{IF } (p_t < ma_t^n(p_t)) \text{ THEN } q_{t+1} = -1 \quad (1.2)$$

The rule starts with a zero investment position. The moving average is calculated at each time step and the quantity q_{t+1} is set to 1 when the price is above its average, -1 when the price is below its average and is unchanged when the price equals its average. The quantity q_{t+1} is the investment position at time $t + 1$ and acts as a multiplier whereby the rule earns q_{t+1} units of the time t to time $t + 1$ return where the return r_t is the log return or log price first difference $r_t = \ln(p_t) - \ln(p_{t-1})$.¹ The investment position is said to be *long* when q_{t+1} is positive, *short* when q_{t+1} is negative and *neutral* when q_{t+1} is zero. Thus, the rule defines a one-step-ahead predictive classification scheme that divides the price into long, short and neutral investment

¹ See Chapter 1 of Campbell *et al.* (1997) for a discussion of returns.

positions that earn positive, negative or zero multiples of the succeeding price change.

A more general version of the price crossover rule is the *moving average crossover rule*. The moving average crossover rule is meant to reduce sensitivity to noise by first smoothing the price. The rule uses two moving averages, a shorter-term average $ma_t^s(p_t)$ of length s and a longer-term average $ma_t^l(p_t)$ of length l where $s < l$. The difference to before is that q_{t+1} now changes when the shorter-term average crosses the longer-term average:

$$q_1 = 0$$

///// RULE

$$q_{t+1} = q_t$$

$$\text{IF } (ma_t^s(p_t) > ma_t^l(p_t)) \text{ THEN } q_{t+1} = 1$$

$$\text{IF } (ma_t^s(p_t) < ma_t^l(p_t)) \text{ THEN } q_{t+1} = -1 \quad (1.3)$$

A more general version still is the *moving average crossover with percentage band rule*. The moving average crossover with percentage band rule is meant to compensate for non-trending dynamics by introducing a percentage band $b \geq 0$ offset around the longer-term average such that the shorter-term average now has to cross the offset before trading takes place. The idea is that in the absence of a trend the shorter-term average tends to wander around repeatedly criss-crossing the longer-term average but where the cumulative price changes between each crossing are insufficient for trading to be profitable (a phenomena known as *whipsawing*). In the presence of a trend however, the shorter-term average is expected to cross the longer-term average and to then continue onwards and cross the offset:

$$q_1 = 0$$

///// RULE

$$q_{t+1} = q_t$$

$$\text{IF } (ma_t^s(p_t) > ma_t^l(p_t)(1 + b)) \text{ THEN } q_{t+1} = 1$$

$$\text{IF } (ma_t^s(p_t) < ma_t^l(p_t)(1 - b)) \text{ THEN } q_{t+1} = -1 \quad (1.4)$$

The moving average crossover with percentage band rule is often used to represent a family of trading rules parameterised as (s, l, b) . The parameters are chosen according to the type of trading rule. For example, $(1, 50, 0)$ defines the 50-day price crossover rule. Note also that the trading rules are always in the market. They do not take neutral positions. A neutral position is said to occur when the price is too close to the moving average to form a view about the trend. To allow for neutral positions, the moving average crossover with percentage band rule can be written as:

$$q_1 = 0$$

///// RULE

$$q_{t+1} = q_t$$

$$\text{IF } (ma_t^s(p_t) > ma_t^l(p_t)(1 + b)) \text{ THEN } q_{t+1} = 1$$

$$\text{IF } (ma_t^s(p_t) < ma_t^l(p_t)(1 - b)) \text{ THEN } q_{t+1} = -1$$

$$\text{IF } (ma_t^s(p_t) < ma_t^l(p_t)(1 + b) \text{ AND } ma_t^s(p_t) > ma_t^l(p_t)(1 - b)) \text{ THEN } q_{t+1} = 0 \quad (1.5)$$

A further refinement is to specify that once a trade is opened, it is held for a fixed period and then closed. Trading rules of this type are defined as having a *fixed length* and are denoted by $FMA(s, l, b)$. For example, if the holding period is fixed at 10 days, once a trade is opened it

is held for 10 days and then closed. Any buy/sell signals that occur after the trade is opened and before it is closed are ignored. If instead the holding period varies according to the time separating the buy/sell signals, the trading rules are defined as having a *variable length* and are denoted by $VMA(s, l, b)$. For example, the trading rules defined by (1.2) to (1.5) are all variable length. The moving average trading rules just described are the most common in the literature. Other moving average trading rules studied in the literature and that are rooted in technical analysis but which receive much less attention are the triple moving average (Wong *et al.* (2003)) and the adaptive moving average (Ellis and Parbery (2005)). Kaufmann (2005) discusses these and other moving average trading rules from a technical analysis perspective.

1.3 Efficient market hypothesis

The distinguishing characteristic common to all technical analysis is that it is conditional on past price information. For this reason, studies of technical analysis usually appear as tests of the weak form *efficient market hypothesis* (EMH). The EMH states that a market is efficient with respect to the information set θ_t if it is not possible to use the information contained in θ_t to formulate buy/sell decisions that earn a higher return than is normal for the same level of risk. A key implication of the EMH is that it is not possible for moving average trading to successfully exploit past prices. This is because markets are *efficient*. An efficient market is defined as a market where the price *fully reflects* all available information. The definition of an efficient market as a market where the price fully reflects all available information is due to Fama (1970) who defines three levels of efficiency, namely, *weak form*, *semi-strong form* and *strong form*, each of which varies according to the information contained in θ_t :

- In a weak form efficient market, θ_t comprises all past price information. In a weak form efficient market, it is not possible to predict future prices using past price information.
- In a semi-strong form efficient market, θ_t comprises all public information. In a semi-strong form efficient market, it is not possible to predict future prices using publicly available information.
- In a strong form efficient market, θ_t comprises all information. This includes price, public and private information.² In a strong form efficient market, it is not possible to predict future prices using any kind of information whatsoever.

Each form of efficiency is progressively less restrictive and includes the information set of its predecessor(s). Thus, if a market is strong form efficient, it is also semi-strong form efficient and weak form efficient. Similarly, if a market is weak form inefficient, it is also semi-strong

² Price information is information relating to the price. Examples of price information are the price itself, whether it is rising or falling and how far it has risen or fallen. Public information is information in the public domain. Examples of public information are news, published end of year company accounts and analyst's earnings forecasts. Private information is information that is known but not by the market. Examples of private information are knowledge of fraud before it is discovered, knowledge of a new order win before it is announced and knowledge acquired through analysis and research. Note that private information can also be interpreted as inside information. Trading inside information is illegal. Inside information is distinct from private information acquired through analysis and research. In market microstructure, private information is used to denote informed trading (see, for example, De Jong and Rindi (2009)). Fama (1991) revises the classification scheme. Tests for weak form efficiency are referred to as tests for return predictability. Tests for semi-strong form efficiency are referred to as event studies. Tests for strong form efficiency are referred to as tests for private information.

form inefficient and strong form inefficient.

A problem with the definition of Fama (1970) is that to test if prices fully reflect all available information it is necessary to test prices against a model that defines precisely what to expect when prices do fully reflect all available information. This is known as the *joint hypothesis problem*. Tests of whether prices fully reflect all available information are joint tests of the hypotheses that (1) the market is efficient and (2) the model against which market efficiency is judged is correct. Tests can fail because one of the two hypotheses is false or because both hypotheses are false.

Jensen (1978) avoids the joint hypothesis problem and stresses the importance of profitability in testing for market efficiency. If it is not possible for a trader to profit financially, evidence of market inefficiency is economically insignificant. Jensen (1978, p. 96) defines the EMH as:

A market is efficient with respect to information set θ_t if it is impossible to make economic profits by trading on the basis of information set θ_t .

Economic profits are defined as risk adjusted returns net of costs. This is a far more practical definition and provides guidance on method. To test for market efficiency, it is sufficient to consider the net risk-return profile of trading rules that trade information set θ_t . This is made explicit in Taylor (2005, p. 175) who defines the EMH as:

No trading rule has an expected, risk adjusted, net return greater than that provided by risk free investment.

The difference between the definition of Fama (1970) and those of Jensen (1978) and Taylor (2005) is that the former does not allow prices to convey predictable information whereas the latter do. While this might seem to be counter intuitive, predictability and profitability are not the same thing. It is possible for prices to be predictable but where the predictable component is not financially exploitable. All definitions are the same in this respect and express the idea that in an efficient market, it should not be possible to systematically outperform the market without taking on excess risk.³

1.4 Moving average trading and market efficiency

The benchmark reference for moving average trading is Brock *et al.* (1992) who test various moving average trading rules applied to the Dow Jones Industrial Average (DJIA) from 1897 to 1986. Ten parameter combinations are evaluated for both fixed and variable length trading rules with $1 \leq s \leq 5$, $50 \leq l \leq 200$ and $b = 0\%$ or 1% . Days when $q_{t+1} = 1$ are classified as *buy* days and days when $q_{t+1} = -1$ are classified as *sell* days. The following results are reported:

- Returns on buy days are consistently higher than returns on sell days. The mean return on buy and sell days across all parameter combinations is 12% and -7% per annum respectively. Tests of the difference in means for buy and sell days are statistically significant for each parameter combination. Results are robust for the sub-periods 1897–1914, 1915–1938, 1939–1962 and 1962–1986. This is inconsistent with the EMH as fully reflecting all available information.

³ See Lim and Brooks (2011) for a systematic review of the empirical literature on the EMH.

- Returns on sell days are negative. The mean return on sell days is negative for all four sub-periods. Returns are not explained by a positive risk premium.
- Returns on buy days are less volatile than returns on sell days. The standard deviation of returns for sell days is less than the standard deviation of returns for buy days for all four sub-periods. If volatility measures risk, the difference in the level of risk does not explain the higher mean return for buy days than for sell days.
- Bootstrap simulations of a random walk with drift, AR(1), GARCH-M and E-GARCH pricing models cannot explain the results. The information uncovered by the trading rules is not explained by linear or time varying volatility models.

Brock *et al.* (1992) view the results as economically significant. The difference in the mean buy – sell spread is 19% per annum compared to a buy and hold return of 5%. They do not allow for costs however. Bessembinder and Chan (1998) test the same trading rules as Brock *et al.* (1992) applied to dividend adjusted DJIA prices from 1926 to 1991. For the full sample period, the mean buy – sell spread across all trading rules is 4.4% per annum, giving one-way breakeven costs of 0.39% per trade. Although significant compared to trading costs estimated at 0.24% – 0.26% per trade, breakeven costs decline over time. For the most recent sub-period of 1976 – 1991, one-way breakeven costs are 0.22% per trade. It is unlikely that traders would be able to use the trading rules to generate profits after costs. Similarly, Hudson *et al.* (1996) test the same trading rules as Brock *et al.* (1992) applied to the FT30 from 1935 to 1994. For the full sample period, the mean round-trip breakeven costs across all trading rules are 0.8% per trade. This is compared to trading costs estimated at 1% upwards. As before, it is unlikely that traders would be able to use the trading rules to generate profits after costs.

1.5 Motivation

The motivation for this thesis is as follows. It is clear from the results of Brock *et al.* (1992) that moving average trading picks up information of some kind. Further, this information is profitable albeit economically insignificant after costs. However, the trading rules are poorly defined. If the buy/sell signals are thought of as defining the space of trading opportunities, the trading rules fail to capture how a trader might react to a trading opportunity in practice. Rather, they simply buy and sell every buy/sell signal regardless of whether this is the right thing to do. The motivation is to address this issue by remodelling the trading rules to better reflect what it is for a trading opportunity to be financially exploitable assuming the resulting trading rules will have more power as a test of the weak form EMH. For example, doubling the breakeven costs of the papers discussed in the previous section would cast doubt on the EMH. The approach is innovative and is to remodel the trading rules to include a description of trading practice. This is fitted to the moving average such that the information that would normally be input to the trading decision is substituted with the information contained in the moving average buy/sell signals. The next section describes the approach to remodelling the trading rules in more detail.

1.6 Remodelling the price crossover rule

Let $buySignal_t$ denote a buy signal and let $sellSignal_t$ denote a sell signal. A general model of the trading rules of Section 1.2 but which does not take neutral positions is:

$$q_1 = 0$$

///// RULE

$$q_{t+1} = q_t$$

IF (*buySignal*_{*t*}) THEN $q_{t+1} = 1$

IF (*sellSignal*_{*t*}) THEN $q_{t+1} = -1$ (1.6)

The question is whether (1.6) can be made to be profitable. By profitable, it is meant that the mean return per trade is greater than the expenditure on costs. Let \bar{R} denote the mean return per trade and let c denote costs. At the end of trading, on average, each trade earns:

$$\bar{R} - c \quad (1.7)$$

This is profitable if and only if:

$$\bar{R} - c > 0 \quad (1.8)$$

The problem with trying to make (1.6) profitable is that to do so with any degree of certainty requires an accurate model of future returns. This is a far from trivial task and is beyond the scope of this thesis for that reason. The approach instead is as follows. As a trend following strategy, a property of moving average trading is that it tends to generate a small number of large *winning* trades offset against a large number of small *losing* trades. A winning trade is defined as one where the return is positive. All other trades are defined as losing trades. In general, a property of the winning trades is that the mean return per winning trade is more than high enough to cover costs. However, there are many more losing trades than winning

trades. The losing trades act to reduce the mean return per winning trade, which reduces the combined mean per trade to the point where it is either not profitable or borderline profitable.

To see this, Table 1.1 lists the mean return per trade for the winning trades, the losing trades and the combined winning and losing trades for various instances of the price crossover rule applied to the daily closing price of Yule Catto (a FTSE Small Cap chemicals company) for the period 03-January-1972 to 30-June-2009. To allow for serial dependence in the order of the trades, the mean return per trade is calculated as the continuously compounded mean. A trade is defined as a sequence of 1's or -1's. Prices include the spread but do not include the dividend. Costs are not included. From this, the mean return per winning trade is more than high enough to cover costs. However, when the winning trades are combined with the losing trades, the combined mean return per trade is loss making. Clearly, the ideal scenario is one where there are winning trades only and no losing trades at all.

Table 1.1
Mean return per trade for Yule Catto

MA	Winning Trades	Winning Mean	Losing Trades	Losing Mean	Combined Mean
10	267	8.854%	762	-3.526%	-0.456%
25	176	10.725%	499	-3.842%	-0.240%
50	110	13.827%	361	-4.068%	-0.158%
100	71	17.373%	272	-4.248%	-0.126%
200	44	21.049%	255	-4.378%	-1.002%

MA is the moving average length n . Winning Trades and Losing Trades are the number of winning trades and losing trades. Winning Mean, Losing Mean and Combined Mean are the mean return per trade for the winning trades, losing trades and combined winning and losing trades.

Let T denote the set of trade that results from applying (1.6) and let this comprise the set of winning trades W and the set of losing trades L such that $W \cap L = \emptyset$ and $T = W \cup L$. The method underpinning this thesis is to use (1.6) to generate T whereby W and L can be said to exist and to then assume that the nature of W and L is consistent with Table 1.1 such that in the limit as the number of losing trades tends to zero, (1.6) is almost surely profitable. The approach is to then remodel (1.6) to keep as many winning trades as possible at the same time as reducing the number of losing trades. This transforms T into $T^* = f(T)$ where $f(T)$ acts on the mean return per trade of T in an attempt to satisfy (1.8). Define \bar{W} and \bar{L} as the mean return per winning trade and the mean return per losing trade:

$$W = \{x \in T \mid x > 0\} \quad (1.9)$$

$$L = \{x \in T \mid x \leq 0\} \quad (1.10)$$

$$\bar{W} = \frac{1}{n_w} \sum_{x \in W} x, \quad n_w = \#\{W\} \quad (1.11)$$

$$\bar{L} = \frac{1}{n_l} \sum_{x \in L} x, \quad n_l = \#\{L\} \quad (1.12)$$

The mean return per trade for both T and T^* is then:

$$\bar{R} = \frac{(n_w \times \bar{W}) + (n_l \times \bar{L})}{n_w + n_l} \quad (1.13)$$

This thesis proposes two new trading rules designed in response to (1.13). The first trading rule is the *trade reduction rule*. The trade reduction rule is motivated by the observation that

when a trade is closed, it is known if it is a winning trade or a losing trade. This information can be used to explicitly transform (1.13) by responding differently to the losing trades. The trade reduction rule follows from the idea of allowing a trade to run. Logically, in overview, the trade reduction rule can be written as:

```
IF (buy / sell signal) THEN
    IF (there is not an open trade) THEN open a trade
    ELSEIF (the trade should be closed) THEN close the trade
    ELSEIF (the trade should be allowed to run) THEN do nothing
END IF
```

(1.14)

The second trading rule is the *positive autocorrelation rule*. The positive autocorrelation rule is motivated by the observation that moving average trading profits from persistence in sign or, equivalently, from persistence in direction. A model that exhibits persistence in sign is to assume returns are positively autocorrelated. By testing for positive autocorrelation, it is less likely that a trade will be a losing trade since it is simultaneously more likely to be a winning trade. This information can be used to implicitly transform (1.13) by responding differently to the buy/sell signals. The positive autocorrelation rule follows from the idea of only trading if it is believed to be profitable to do so. Logically, in overview, the trading rule can be written as:

```

IF (buy / sell signal) THEN
    IF (there is an open trade) THEN close the trade
    ELSEIF (opening a trade is likely to be profitable) THEN open a trade
    ELSEIF (opening a trade is unlikely to be profitable) THEN do nothing
END IF

```

(1.15)

Both trading rules extend the price crossover rule to include a *do nothing* response. Due to its design, the trade reduction rule is limited to remodelling the price crossover rule. It does not generalise to allow for the other types of trading rule defined in Section 1.2. This is because the buy/sell signals generated by the price crossover rule are ordered as minima and maxima and this property is exploited by the trade reduction rule. There is no equivalent for the other types of trading rule. The positive autocorrelation rule does generalise although it is difficult to imagine what value there is in this. Consequently, the positive autocorrelation rule is also limited to remodelling the price crossover rule. Both trading rules are exposed to exactly the same buy/sell signals as the price crossover rule. The difference between the trade reduction and positive autocorrelation rules is that the trade reduction rule applies *after* a trade is open whereas the positive autocorrelation rule applies *before* a trade is open. The trading rules can therefore be combined to give both sides of the trading process.

1.7 Research objectives

The research objectives are threefold. First, they are to build the proposed trading rules. The trading rules are straightforward, easy to replicate and are tied to the price crossover rule via the buy/sell signal generating mechanism. They do not deviate from moving average trading to the point where it is no longer recognisable as such. Second, they are to test if the trading

rules are profitable and if they improve the price crossover rule. Third, they are to determine whether the trading rules are economically significant.

1.8 Scope

Inevitably with work on trading systems, there is the question of method. It must be stressed that the thesis is limited to the technical analysis method of moving average trading only and that methods/strategies/technologies outside this are not considered. There is also no attempt to build a working system. Related problems such as how to choose which stocks to trade and how to choose between the trading rules are not addressed empirically. The scope is limited to testing the trading rules as defined by the research objectives only. Testing is on a trading rule by trading rule basis. There is no optimisation.

1.9 Contributions

There are three contributions. First, there are the trading rules. Both trading rules have been designed to capture trading practice and include elements of decision-making as found in the real world. Second, there is evidence that the trading rules uncover information that is missed by the price crossover rule and that this information is financially exploitable. This motivates the argument that the information needed for trading to be economically viable is observable in the price. This is a challenge to market efficiency as defined by Fama (1970). Third, there is an empirical link with market microstructure. The trading rules uncover issues of informed trading (asymmetric information), liquidity, adverse selection and price impact. The strongest interpretation that can be applied to the trading rules in this context is that they are examples of informed trading. It is not known of any work that explicitly links moving average trading and market microstructure in this way.

1.10 Thesis outline

The remainder of the thesis is organised as follows. Chapter 2 reviews the literature. Chapter 3 describes the data and test method used to test the trade reduction rule. Chapter 4 presents the trade reduction rule. Chapter 5 presents the positive autocorrelation rule. Chapter 6 offers conclusions and provides direction on further work.

Chapter 2

Literature Review

This chapter reviews the literature on technical analysis. Particular attention is paid to moving average trading. The chapter starts by defining technical analysis and provides an overview of the main techniques. These are charts, trading rules and cycle analysis. This is followed by a review of the survey literature on the use of technical analysis by professional traders in the foreign exchange and equity markets and which is used to motivate the point that the trading rules studied in the literature bear little resemblance to trading practice. The remainder of the chapter is organised as follows. Section 2.1 introduces technical analysis and summarises the main techniques. Section 2.2 reviews the survey literature. Section 2.3 offers conclusions.

2.1 Technical analysis

Technical analysis is an approach to predicting future price movements based on identifying patterns in prices, volume and other market statistics. The philosophy underpinning technical analysis is that future prices are predictable from past prices as long as prices reflect changes in supply and demand. The approach is to detect trends as soon as possible and to trade in the trend direction. Pring (2002, p.2) describes technical analysis as:

The technical approach to investment is essentially a reflection of the idea that prices move in trends that are determined by changing attitudes of investors

towards a variety of economic, monetary, political and psychological forces. The art of technical analysis, for it is an art, is to identify a trend reversal at a relatively early stage and ride on that trend until the weight of the evidence shows or proves that the trend has reversed.

Technical analysis differs from fundamental analysis. Fundamental analysis uses economic variables such as interest rates, valuation ratios and industry trends to predict future returns based on an economic model. Technical analysis uses past prices and other measures of the price to predict future returns based on extrapolating the price into the future. Fundamental analysis is usually concerned with predicting the long term and guides *investment* decisions. In contrast, technical analysis is usually concerned with predicting the short term and guides *trading* decisions. Fundamental analysis and technical analysis have little in common in this respect. See Welch and Goyal (2008) for a survey of return predictability based on economic variables.

To this end, technical analysis employs a number of techniques, the most common of which are *charts*, *trading rules* and *cycle analysis*. Charting relies on detecting graphical patterns in the price. Patterns are usually defined as reversal and continuation patterns. Reversal patterns include the head and shoulders, double tops/bottoms and rounded tops/bottoms. Continuation patterns include flags, pennants, wedges and rectangles. Studies of charting are often limited by the need to design a pattern recognition algorithm to extract the patterns although studies of charting are becoming increasingly common. See, for example, Lo *et al.* (2000), Dempster and Jones (2002), Dawson and Steeley (2003), Wang and Chan (2007, 2009) and Leigh *et al.* (2008). The general result is that there is evidence of predictive ability. It is not clear to what extent this equates to profitability however.

Because they are mathematically tractable, the majority of studies of technical analysis are presented in the form of trading rules. As defined in Chapter 1, a trading rule is a numerical method that maps the price onto investment decisions. Trading rules are based on *technical indicators* where a technical indicator is a quantitative function of the price state. Technical indicators include moving averages, momentum oscillators and volume indicators. Example momentum oscillators are the %K stochastic and the relative strength index (RSI). Example volume indicators are the on balance volume and the money flow index. See Achelis (2001) for an overview of the more popular technical indicators. Trading rule studies include Mills (1997), Sullivan *et al.* (1999), Day and Wang (2002), Kwon and Kish (2002), Olson (2004), Marshall and Cahan (2005) and Marshall *et al.* (2009). The general result is that while there is evidence of predictive ability, the trading rules are rarely profitable after costs.

Cycle analysis decomposes the price into cycles or trends with different frequencies. Cycle analysis includes Dow Theory, Elliot Wave Theory and Kitchin Waves. It also worth noting that cycle analysis includes seasonality's such as the January effect (Atanasova and Hudson (2010)). Cycle analysis is the least represented in the literature due to its esoteric nature. For example, Dow Theory decomposes the price into a primary trend, secondary trend and minor trend. Primary trends are further decomposed into phases (Achelis (2001)). An exception is Brown *et al.* (1998) who find that Dow Theory as defined by the recommendations made by William Peter Hamilton during his tenure as editor of the *Wall Street Journal* for the period 1902 – 1929 results in positive risk adjusted returns. This is before costs however.

Park and Irwin (2007) review the profitability of technical analysis and discuss charting and trading rule approaches as well as approaches based on genetic algorithms, neural networks

and time series models. The conclusion is that despite evidence of predictability, improved testing procedures and plausible theoretical explanations, academic acceptance of technical analysis is hindered by problems of data snooping and potentially economically insignificant profits after adjusting for risk. This is attributed to a lack of understanding of how technical analysis is used in real world markets.

2.2 Survey literature

Historically, technical analysis has been treated with scepticism in the academic literature. This can be linked to (1) influential and widely cited early studies of technical analysis such as Alexander (1961, 1964), Fama and Blume (1966), Van Horne and Parker (1967, 1968) and Jensen and Benington (1970) all of whose findings are negative and which do not support technical analysis as having predictive value, (2) the dominance of the EMH as the prevailing theoretical paradigm resulting in proponents of the EMH such as Malkiel (1996) dismissing technical analysis as worthless, (3) the fact that much of technical analysis lacks a strictly logical explanation and (4) the lack of evidence from practitioners of technical analysis to support any claims to the contrary.

Nevertheless, technical analysis remains widely used. Survey studies find strong evidence of the continued use of technical analysis in practice. Taylor and Allen (1992) survey the use of technical analysis amongst chief foreign exchange dealers in London in 1988. They find that at least 90% of respondents place some weight on technical analysis in their decisions. The weight given to technical analysis is highest for short forecast horizons of intra-day to three months and declines for longer forecast horizons of six months to a year, where greater weight is given to fundamental analysis. The results indicate that 64% of respondents use moving averages and/or trend following systems and that 40% use rate of change indicators and/or

oscillators. Technical analysis and fundamental analysis are seen as complementary with 60% of respondents considering technical analysis at least as important as fundamental analysis. Anecdotal evidence implies that technical analysis is used to confirm fundamental analysis rather than to contradict it. This suggests that technical analysis is used mainly in a decision support capacity and not as a trading vehicle in its own right. However, 2% of respondents rely exclusively on technical analysis and do not appear to use fundamental analysis at all.

Menkhoff (1997) surveys foreign exchange dealers and fund managers in Germany in 1992. Technical analysis is used extensively with 87% of respondents giving at least a 10% weight to technical analysis in their decisions. As in Taylor and Allen (1992), technical analysis and fundamental analysis are seen as complementary. The weight given to technical analysis does not decline so rapidly as the forecast horizon increases however and is relatively constant for all forecast horizons. Menkhoff (1997) also surveys the use of flow analysis (the information content in the order flow).¹ The weight given to flow analysis is highest for forecast horizons of intra-day and declines as the forecast horizon increases. The mean weights attached to the importance of flow analysis, technical analysis and fundamental analysis are 18%, 37% and 45% respectively. Users of technical analysis are not differentiated by seniority or education. There is a slight preference for technical analysis among younger respondents but this is not statistically significant.

Lui and Mole (1998) survey foreign exchange dealers in Hong Kong in 1995. Questions on how technical analysis is used are very specific. Technical analysis and fundamental analysis

¹ Flow analysis is the analysis of information contained in the order flow where order flow is signed transaction volume indicating purchases and sales. See Gehrig and Menkhoff (2004).

are seen as complementary with 85% of respondents confirming their use. The weight given to technical analysis is highest for short forecast horizons and declines as the forecast horizon increases, where greater weight is attached to fundamental analysis. Fundamental analysis is seen as more useful in predicting trends whereas technical analysis is seen as more useful in predicting turning points. The most popular technical analysis methods are moving averages and/or trend following systems. The typical length of historic data is 12 months. Daily data is the most popular.

Oberlechner (2001) surveys foreign exchange traders and financial journalists in Frankfurt, London, Vienna and Zurich in 1996. Although difficult to interpret, for both groups, greater weight is given to technical analysis for short forecast horizons and decreases as the forecast horizon increases, where greater weight is given to fundamental analysis. Traders attach the highest weight to technical analysis for forecast horizons of up to three months. Journalists attach the highest weight to technical analysis for forecast horizons of up to one month. The weight given to fundamental analysis by journalists is always higher than that of traders. This suggests that those inside the market see it differently from those outside the market although this can also be explained by a journalistic need to rationalise events. The majority of traders use both technical analysis and fundamental analysis with 3% relying exclusively on one of the two methods. When the results for traders in London are compared with Taylor and Allen (1992), the importance attached to technical analysis increases across all forecast horizons. Users of technical analysis are not differentiated by age, gender, market type or seniority. However, traders with a less than US\$50 million limit demonstrate a statistically significant preference for technical analysis compared to traders with a greater than US\$50 million limit.

Cheung and Chinn (2001) survey foreign exchange traders in the United States between 1996 and 1997. Respondents were asked to *best describe* their trading practice. Technical trading best describes 30% of trading practice. Fundamental analysis best describes 25% of trading practice. The remainder is characterised as customer order driven (22%) and jobbing (23%). The percentage of respondents who best describe their trading practice as technical trading also appears to have increased. When asked to best describe their trading practice five years ago, 19% of respondents describe themselves as technical traders. The increase in technical trading is at the expense of jobbing. The percentage of respondents who best describe their trading practice five years ago as fundamental analysis is 23%.

Similarly, Cheung *et al.* (2004) survey foreign exchange traders in the United Kingdom in 1998. As in Cheung and Chinn (2001), respondents were asked to best describe their trading practice. Technical trading best describes 33% of trading practice. Fundamental analysis best describes 34% of trading practice. The remainder is characterised as customer order driven (37%) and jobbing (36%). When asked to best describe their trading practice five years ago, 14% of respondents describe themselves as technical traders. The increase in technical trading is at the expense of jobbing. The percentage of respondents who best describe their trading practice five years ago as fundamental analysis is 31%.

Gehrig and Menkhoff (2006) survey foreign exchange dealers and fund managers in Austria and Germany in 2001. The authors compare the results to Menkhoff (1997). Overall, the use of technical analysis has gained ground. Technical analysis is used extensively with 97% of respondents giving at least a 10% weight to technical analysis compared to 87% before. The mean weights given to the importance of flow analysis, technical analysis and fundamental analysis by dealers are 26%, 42% and 32%. This is compared to 21%, 37% and 42% before.

For fund managers, the mean weights are 17%, 37% and 46%. This is compared to 9%, 37% and 54% before. The importance of flow analysis and technical analysis has increased while the importance of fundamental analysis has declined. Users of technical analysis are also not differentiated by age, seniority, company size or education.

The survey studies discussed so far concentrate on the foreign exchange. The only known survey study of the equity markets is Menkhoff (2010) who surveys mutual, pension, bond and equity fund managers in the United States, Germany, Switzerland, Italy and Thailand in 2003/2004. Technical analysis is used extensively with 55% – 87% of respondents giving at least a 10% weight to technical analysis in their decisions. The mean weights given to flow analysis, technical analysis and fundamental analysis are 10%, 23% and 67%. The greatest weight is given to flow analysis for forecast horizons of intra-day to days, technical analysis for forecast horizons of weeks and fundamental analysis for forecast horizons of months to years. An average of 20% of respondents prefer technical analysis relative to other forms of analysis. This is attributed to the use strategies that rely to a large degree on technical input such as momentum.

2.3 Conclusion

The main point to take from the literature review is the simplest one – the trading process is information rich and highly complex. The different weights given to flow analysis, technical analysis and fundamental analysis suggest different classes of information, all of which act as input to the trading decision. Further, the variation in forecast horizon implies that traders are likely to fit this information to a number of different models that are not only specific to the type of trade but also to the perceived implications of the information content. Given the survey literature, to dismiss technical analysis on the back of the argument that a trading rule

such as the moving average trading rules defined in Section 1.2 is sufficiently representative of trader behaviour is highly questionable. It is by understanding traders that it is possible to understand technical analysis and not the reverse.

To develop this point, the trading rules proposed in this thesis are designed to capture trading practice. Specifically, they are designed to capture the trading process as driven by *direction*. That is, the idea of prices going up and down and trading as an attempt to profit from this. To achieve this, it is necessary to model the price as having direction. This is modelled using the moving average. The trading rules extend moving average trading to include a description of trading practice and are motivated by the desire to explore the nature of trading as typified by the survey papers just discussed. The remainder of the thesis expands on the trading rules in more detail.

Chapter 3

Data and Test Method

This chapter describes the data and test method used to analyse the trade reduction rule in the next chapter. The data comprises 45 years of daily close and bid-ask prices for stocks listed on the London Stock Exchange from 01-January-1965 to 30-June-2009. The stocks are drawn from the FTSE 100, FTSE 250, FTSE Small Cap and FTSE Fledgling indices. An issue with the data is that the bid-ask prices are often inconsistent with the close price in that they do not always conform to the ordering $bid_t \leq close_t \leq ask_t$. This means it is necessary to recreate the spread despite the availability of bid-ask prices and the majority of this chapter is directed at addressing this issue. The test method takes the form of hypothesis tests. All hypothesis tests are bootstrap tests. The remainder of the chapter is organised as follows. Section 3.1 describes the data and the approach to recreating the spread. Section 3.2 provides an outline of relevant simulation issues. Section 3.3 defines the hypothesis tests. Section 3.4 concludes with a short summary.

3.1 Data

This section describes the data. Section 3.1.1 discusses the requirements of the data. Section 3.1.2 describes the data content and format. Section 3.1.3 describes the pre-processing applied to the close price. After the close price has been pre-processed, the data is considered useable. However, a general issue with the data is that the bid-ask prices are often inconsistent with the

close price. Section 3.1.4 states that the reason for this is the sampling of bid-ask prices after the mandatory quote period and describes the approach to recreating the spread when dealing with inconsistent or missing bid-ask prices. Section 3.1.5 describes the breakdown of the data into portfolios and sub-periods.

3.1.1 Requirements and considerations

The requirements of the data are twofold. Firstly, it is to provide a data set by which to test the trading rules and, secondly, it is to provide a measure of cost. With reference to the latter requirement, an issue with moving average trading is that in the absence of costs the trading rules overstate the return. This is because the trading rules assume trending behaviour in the underlying price dynamics but when the price does not trend the trading rules are sensitive to noise. This can sometimes generate small positive returns of, say, 1% but which are negative after costs. If costs are not allowed for, some trades are misclassified as winning trades when they are in fact losing trades. This overstates the return and can lead to erroneous conclusions regarding profitability. Allowing for costs militates against this.

Costs in this case equate to the spread.¹ Note however that the spread is subject to estimation error. There are two reasons for this. Firstly, if available, the spread is not guaranteed to be error free and, secondly, there are a large number of prices where the spread is not available. Both cases require the spread to be estimated resulting in estimation error. Underlying this is the question of experiment design. Given that prices are available from 1965 to 2009, which is a large sample, it appears sensible to test all of them. However, the spread is not available until 1986 at the earliest. If prices before 1986 are not tested on the basis that the spread is not

¹ The spread is an execution cost and is the only cost considered. See Section 3.1.4.

available before this, half the data is abandoned. If prices before 1986 are tested on the basis that it is possible to estimate the spread before this, half the data is subject to estimation error. This has to be balanced against questions of do the trading rules work, are the results robust in the presence of costs and are the results robust over time.

Because the concern is to test if the trading rules work, which raises the issue of whether the trading rules work for all prices for all time, the approach is to test all of the data with prices before 1986 seen as a robustness check. Consequently, the data is split into three sub-periods of fifteen years each with 1986 as the middle of the middle sub-period. This gives three data sets where the change in the availability of the spread is phased as opposed to stepped. Note however that it is still necessary to estimate the spread due to it not being error free. It is of course possible to avoid this by using different data. Given that it is now possible to obtain high frequency time-stamped records of daily order flow and quote revisions, this data would eliminate a lot of error. However, the data sets are large and require significant filtering (see, for example, Dacorogna *et al.* (2001)). This is a considerable undertaking and is unnecessary at this stage. The requirement is not to analyse the spread in itself, but is instead to derive a reasonable estimate of costs. The test data is sufficient for this purpose.

3.1.2 Content and format

The data consists of 45 years of daily close and bid-ask prices for stocks listed on the London Stock Exchange from 01-January-1965 to 30-June-2009. The data is supplied by DataStream and comes as a sequence of date-stamped records of the form $(date_t, bid_t, close_t, ask_t)$ where $date_t$ is the date for which prices are available, bid_t is the end of day bid price, $close_t$ is the end of day close price and ask_t is the end of day ask price. Dividends are not included. Prices

are at daily intervals and exclude weekends but include public holidays. The exchange closes on public holidays and prices for public holidays are duplicates of the previous days prices. Prices for public holidays are not treated differently and are retained along with other prices. The close price is always available and starts from 01-January-1965. The bid-ask prices are available from 27-October-1986 onwards following the so-called “Big Bang”. The bid-ask prices do not always start from this date though, 27-October-1986 is the earliest date from which they are available. It is common to find the close price starting before this date and for the bid-ask prices to start after this date. There is one price file per stock and there are 567 stocks in total. Investment trusts are not included.

The data is available in two different formats. The master price in both formats is the close price. The first format is unadjusted and is the official close price quoted on the exchange. The second format is adjusted for capital events such as stock splits and rights issues. A historical adjustment factor that accumulates capital events in reverse chronological order back to the base date is also available. In general, it is easier to work with adjusted prices since there is no need to manage changes in the price level attributable to changes in capital structure. Prices are adjusted prices for this reason. Returns are not affected but it does mean that some of the prices have never existed.

3.1.3 Pre-processing the close

The first step in pre-processing is to check that there are no missing close prices. All trading rules use the close price to generate buy/sell signals and this can cause problems if the close price is missing. The first check is to delete all price files where the close price is missing. In general, DataStream is a respected data vendor and it is unusual to find missing close prices. The main reason for missing close prices is when a stock is suspended. The next check is to

test whether the adjusted $(bid_t, close_t, ask_t)$ prices equal the unadjusted $(bid_t, close_t, ask_t)$ prices multiplied by the adjustment factor. Testing is to zero decimal places. If the adjusted and unadjusted prices match, the price file is accepted. If they do not match, the price file is inspected manually. If it is clear why the test fails, the error is fixed and the test is repeated. Examples of why a test fails are rounding errors and a delay in the update of the adjustment factor. If it is not clear why a test fails or the error persists in some way, the price file is deleted.

3.1.4 Reconstructing the spread

The next step in pre-processing is to check whether the $(bid_t, close_t, ask_t)$ prices conform to the ordering $bid_t \leq close_t \leq ask_t$. When trading takes place it takes place at one of two prices, namely, the *bid* price or the *ask* price. The bid price is the price received when selling and the ask price is the price paid when buying. The difference is the *bid-ask spread*. In percent, the bid-ask spread $s(ba)_t$ at time t is:

$$s(ba)_t = \frac{ask_t - bid_t}{mid_t} \times 100\%, \quad mid_t = \frac{bid_t + ask_t}{2} \quad (3.1)$$

The spread defines the costs incurred for each round-trip transaction and is the cost paid by the trading rules for being able to trade immediately at the prevailing price. Equation (3.1) is also known as the *quoted spread*. It should be mentioned that there is a substantial literature dedicated to analysing the spread wherein the spread arises to compensate market makers or liquidity providers for transitory (non-informational) and adverse selection (informational) costs. Bessembinder and Venkataraman (2010) discuss this literature in more detail. For the purposes of the thesis however, the spread is a cost levied against the trading rules assuming

all trades are executed as market orders. The trading rules demand liquidity and the spread is the price paid for it.

There are two problems with the spread. The first problem is that the bid-ask prices are only available from 27-October-1986 onwards. This is not overly critical but it does mean that it is necessary to reconstruct the spread before this. The second problem is that the bid-ask prices are often inconsistent with the close price. The close price is always the official close price and is always the master price. The problem is that the bid-ask prices appear to be sampled after the mandatory quote period ends at 16:30 GMT. The mandatory quote period runs from 08:00 to 16:30 GMT and market makers are obliged to offer firm two-way prices during this time. Outside the mandatory quote period, market makers are not obliged to offer prices if they do not wish to do so. This leads to greater price uncertainty and the bid-ask prices can move away from the close price because of this. A straightforward test of the consistency of the bid-ask prices is to check if they conform to the ordering $bid_t \leq close_t \leq ask_t$. All price files fail this test at least once. This means that it is necessary to recreate the spread despite the availability of the bid-ask prices for a particular date. However, it can be assumed that errors in the bid-ask prices are due to sampling after the mandatory quote period. A more general consequence of sampling the bid-ask prices after the mandatory quote period is that the spread is likely to be overestimated. In a study of intra-day data for example, Abhyankar *et al.* (1997) find that the spread widens up to 30% after the mandatory quote period ends at 16.30 GMT. In a similar study, Cai *et al.* (2004) exclude data outside the mandatory quote period on the basis that it is indicative only.

The first step in reconstructing the spread is to zero each bid price greater than the close price and each ask price less than the close price. Zeroing is on an individual basis. If the bid price

is greater than the close price and the ask price is not less than the close price for example, the bid price is set to zero and the ask price is left unchanged. The next step is to zero all bid-ask prices that equal one another. Each price entry then conforms to either $bid_t \leq close_t < ask_t$ or $bid_t < close_t \leq ask_t$ but not both. Because the bid-ask prices are sampled after the mandatory quote period, the next step is to identify bid-ask prices that have moved so far away from the close price that they appear as outliers. The problem is that it is difficult to define precisely what an outlier is. For example, it is perfectly reasonable for a stock to trade with a spread of 1%, for the business to fail in some way, for the share price to crash and for the stock to end up trading as a penny share with a spread of 100%. A spread of 100% is then unusual at the start of the price file but not at the end. This suggests a method based on moving windows. To expedite this process, the method is to adopt a “you’ll know it when you see it” approach and to manually inspect the bid-ask prices of each price file for deviation from their nearest neighbours whenever the spread is greater than five times the mean. The choice of five times the mean is firstly to ensure that bid-ask prices less than this are not filtered and, secondly, to identify bid-ask prices that are reasonably distant from the mean. The mean is preferred over the median since it is less robust to outliers and is less likely to result in bid-ask prices being removed.

Because either the bid price, the ask price or both the bid and ask price can be in error, the method is to compute separate bid-ask spreads for the non-missing bid-ask prices. This is always relative to the close price. The close price is then used as the reference price with which to reconstruct the spread. The bid spread $s(b)_t$ relative to the close price at time t is:

$$s(b)_t = \frac{close_t - bid_t}{close_t} \tag{3.2}$$

The ask spread $s(a)_t$ relative to the close price at time t is:

$$s(a)_t = \frac{ask_t - close_t}{close_t} \quad (3.3)$$

Bid-ask prices greater than five times the mean are inspected manually. If the bid-ask price is inconsistent with its nearest neighbours and can be attributed to sampling after the mandatory quote period, it is set to zero, the mean is recalculated and the process is repeated. Zeroing is on an individual basis. The process repeats until either there are no bid-ask prices greater than five times the mean or, if there are, where the bid-ask prices are either consistent with their nearest neighbours or cannot be attributed to sampling after the mandatory quote period.

The next step is to go through each price file and to set each zeroed bid-ask price to have the same spread as its immediate predecessor. Bid-ask prices are reset on an individual basis. The objective is to preserve local asymmetry in the spread. Zeroed bid prices are estimated by:

$$bid_t = close_t \left(1 - \frac{close_{t-1} - bid_{t-1}}{close_{t-1}} \right) \quad (3.4)$$

Zeroed ask prices are estimated by:

$$ask_t = close_t \left(1 + \frac{ask_{t-1} - close_{t-1}}{close_{t-1}} \right) \quad (3.5)$$

Note that as stated in Section 3.1.1, the requirement is to derive a reasonable estimate of the costs incurred by the trading rules given that they will otherwise overstate the return. It is not to study the spread in itself. The straightforward nature of (3.4) and (3.5) reflect this. Methods such as, say, differential equations or Roll's (1984) estimator introduce a level of complexity that is difficult to justify assuming that if the results are significant in any way, the next step would be to turn to the high frequency data at which point the trading rules can be simulated with greater accuracy and in real time.

The final step is to estimate missing bid-ask prices back to the base date. Whatever method is used will introduce a bias since the spread is unknown before this. The method is to calculate the mean of (3.2) and (3.3) using the first available 260 bid-ask prices where a bid-ask price is only included in the calculation if its spread is different to its immediate predecessor. This is to guard against reusing previously missing bid-ask prices as well the situation where prices do not change for long periods, as is sometimes the case with smaller stocks. If there is not enough data to calculate the mean using 260 entries, all the bid-ask prices are used provided their spread is different to their immediate predecessors. Let $\overline{s(b)}$ denote the resulting mean bid spread. Missing bid prices back to the base date are estimated by:

$$bid_t = close_t (1 - \overline{s(b)}) \quad (3.6)$$

Let $\overline{s(a)}$ denote the resulting mean ask spread. Missing ask prices back to the base date are estimated by:

$$ask_t = close_t (1 + \overline{s(a)}) \quad (3.7)$$

Table 3.1 lists the mean bid-ask spread defined in (3.1). The bid-ask spread is calculated after the spread has been reconstructed and where the price files have been grouped into portfolios and sub-periods. The grouping of the price files into portfolios and sub-periods is described in the next section. The mean bid-ask spread is the mean of all the price files for each portfolio for each sub-period. The spreads for FTSE AIM 100 and FTSE AIM All Share portfolios are also shown for comparison. The AIM portfolios are not part of the test data however.

Table 3.1
Mean bid-ask spreads

Period	FTSE 100	FTSE 250	FTSE Small Cap	FTSE Fledgling	FTSE AIM 100	FTSE AIM All Share
1	1.399%	2.116%	3.570%	6.262%	2.495%	5.952%
2	1.422%	2.513%	4.586%	6.990%	3.685%	7.155%
3	0.708%	1.712%	3.263%	7.011%	4.710%	11.882%
All	1.077%	2.006%	3.699%	6.879%	4.323%	11.044%

Period 1 refers to the test period 01-Jan-1965 to 31-Dec-1979. Period 2 refers to the test period 01-Jan-1980 to 31-Dec-1994. Period 3 refers to the test period 01-Jan-1995 to 30-Jun-2009. Period All refers to the test period 01-Jan-1965 to 30-Jun-2009. Variances are not calculated since the spread is bounded below by zero and hence the distributions are likely to be right skewed implying negative spreads.

The spread increases as company size decreases and is lowest for the FTSE 100 and highest for the FTSE AIM All Share. The FTSE AIM 100 also trades on a tighter spread than might be expected. Figure 3.1 plots the mean annual bid-ask spread. The spread before 1990 is less variable than the spread after 1990 although, in general, it does not appear to be significantly distorted. However, it does not capture the overall variability of the spread and it is unlikely that the bid-ask prices would have been achievable in practice. It is also worth noting that in October 1997 the London Stock Exchange introduced a limit order book for the FTSE 100

and more liquid members of the FTSE 250 via the Stock Exchange Trading System (SETS). The limit order book allows public traders (non-market makers) to post prices at which they are prepared to trade and hence to act as counterparties in the supply and demand of liquidity. A limit order book was subsequently introduced for non-SETS members of the FTSE 250 in November 2003 via SETSmm (SETS with market makers), which was extended to cover the FTSE Small Cap and FTSE AIM 50 indices in July 2005 and December 2005 respectively. It is noticeable from Table 3.1 that this appears to have resulted in a decrease in the spread from test period 2 to test period 3 and is indicative of a more cost efficient trading mechanism. A similar decrease is also evident in Figure 3.1.




Figure 3.1 Mean annual bid-ask spreads.

In general, the decrease in the spread would also tend to imply an increase in *liquidity*. It is common in the literature to use the spread as a proxy for liquidity where liquidity is defined as the ability to trade quickly at low cost and with minimal price impact and where liquidity decreases as the spread increases (alternatively, illiquidity increases as the spread increases).

There are problems with the spread as a proxy for liquidity however. First, the spread has to be available. If not, it is necessary to rely on other measures. Lesmond (2005) discusses this issue for emerging markets. Second, for daily spreads, the spread is likely to refer to the last trade of the day only, which may or may not be representative of liquidity during the course of the day. Spreads obtained at daily frequency are noisy in this respect. Third, the spread is likely to reflect the detail of the underlying market structure. Cai *et al.* (2008) find that the spread decreases for SETS traded stocks following its introduction in October 1997, which suggests that the spread is likely to be wider for quote driven markets than for order driven markets. Fourth, the spread is open to manipulation by market makers, a famous example of which is the “NASDAQ controversy” (Christie and Schultz (1994), Christie *et al.* (1994)). Fifth, the spread acts an invitation to trade. It does not convey information about whether a trade will take place inside or outside the spread. The spread conceals the depth of liquidity. Last, while the spread is at least indicative of liquidity in that it conveys information on the cost of liquidity, it does not convey information on the difficulty of executing a trade or the magnitude of its price impact. This is typical of the view of liquidity as a multi-dimensional variable that incorporates *speed* (time taken to execute a trade), *tightness* (low trading costs), *depth* (the volume possible without affecting the price), *breadth* (the number of participants actively engaged in the market) and *resilience* (the speed at which price fluctuations due to the trading process die out) (Hasbrouck (2007)). Goyenko *et al.* (2009) discuss liquidity and its measurement in more detail.

Given the one-dimensional nature of the spread, there are also issues regarding its ability to predict liquidity in the long term. First, the spread is asset specific. If liquidity is driven by factors exogenous to the spread, whether the spread is sufficient to characterise liquidity is debatable and there is now a growing literature that studies the determinants of liquidity by

testing for co-movement or commonality across alternative measures of liquidity (see, for example, Chordia *et al.* (2000), Korajczyk and Sadka (2008) and Brockman *et al.* (2009)).

Second, liquidity is dynamic. For example, suppose there is a perceived increase in risk and that this causes liquidity to flow from illiquid stocks to liquid stocks, which in turn causes the spread for the illiquid stocks to increase and the spread for the liquid stocks to decrease. As a minimum, this suggests that the information content in the spreads of assets of the same type is potentially asymmetric, functionally interdependent and may exhibit a lead-lag relationship.

Last, markets are linked. If liquidity flows out of one market or asset class, it is reasonable to expect it to flow into another. In other words, the relationship between measures of liquidity across different types of market and asset class might be a better predictor of future liquidity compared to measures of the market or asset class on their own. Chordia *et al.* (2005) and Goyenko and Ukhov (2009) find that the liquidity of US bond and stock markets are linked with a change in the liquidity of one market affecting liquidity in the other. The authors also find evidence that the bond market acts as a channel for the transmission of monetary policy variables via changes in liquidity. Overall, the preceding discussion implies that predicting liquidity is non-trivial and is likely to involve economic and microstructure variables, either directly or indirectly, as well as some function of their variation caused by changes in capital movements.

3.1.5 Portfolios and sub-periods

Once the data is pre-processed, it is divided into portfolios. Portfolios are differentiated by company size. The portfolios comprise stocks drawn from the FTSE 100, FTSE 250, FTSE Small Cap and FTSE Fledgling indices. The FTSE 100 index comprises the 100 most highly capitalised companies listed on the main market and represents approximately 80% of UK market capitalisation. The FTSE 250 index comprises the next 250 most highly capitalised

companies and represents approximately 15% of UK market capitalisation. The FTSE Small Cap index comprises the remaining companies listed as members of the FTSE All Share and represents approximately 2% of UK market capitalisation. The FTSE All Share index is the aggregation of the FTSE 100, FTSE 250 and FTSE Small Cap indices. The FTSE Fledgling index comprises companies listed on the main market but which are too small to be included in the FTSE All Share index.

The classification scheme is that if a stock is a member of the FTSE 100 index, it is allocated to the FTSE 100 portfolio. If it is a member of the FTSE 250 index, it is allocated to the FTSE 250 portfolio and so on. The indices are supplied by the FTSE Group (www.ftse.com) and are correct as of August 2009. It should be noted that the FTSE indices are revised on a quarterly basis and are updated to reflect changes in market capitalisation. The portfolios do not reflect these changes. Instead, they reflect the indices as of August 2009 and do not vary with time. The reason for this is practical. An issue with moving average trading is that it suffers from a stock selection problem. If the price dynamics exploited by moving average trading are not in the price, the trading rules will not find them. Therefore, for the trading rules to have practical value, it is necessary to solve the stock selection problem. The stock selection problem in this case is seen as a search problem. Given an investment universe, the stock selection problem is to search the investment universe and to select those stocks most suited to the trading rules. A potential solution to this problem is proposed in Chapter 5. The classification scheme mirrors this perspective in that it splits the investment universe into pseudo arbitrary search spaces at the same time as providing guidance on where to search. While it can be argued that a ranking scheme that explicitly controls for size also addresses this issue and simplifies the analysis in general, the decision not to use a ranking scheme is attributable to the proposed solution to the stock selection problem. That said however, an investigation of the stock selection problem is

further work and there is no reason why a ranking scheme or cross-sectional study cannot be employed for intermediate analytical purposes.

The portfolios, then, are limited to the main market and comprise 99 stocks for the FTSE 100 portfolio, 197 stocks for the FTSE 250 portfolio, 159 stocks for the FTSE Small Cap portfolio and 112 stocks for the FTSE Fledgling portfolio. A full listing of each portfolio including start and end dates is given in the appendix. AIM stocks are not included on the basis that the AIM constitutes a different market although this is not to say that the trading rules cannot be tested on the AIM because of this.

The data is also divided into sub-periods of 15 years each. This is to allow for robustness and to test whether the returns to the trading rules vary with time. The first sub-period is referred to as *test period 1* and is from 01-January 1965 to 31-December-1979. The second sub-period is referred to as *test period 2* and is from 01-January 1980 to 31-December-1994. The third sub-period is referred to as *test period 3* and is from 01-January 1995 to 30-June-2009. The last period is referred to as *test period All* and covers all of the data from 01-January 1965 to 30-June-2009. Table 3.2 summarises the breakdown into portfolios and test periods in terms of the number of stocks in each portfolio for each test period. All test periods are reasonably well populated. Understandably, the number of stocks in each test period increases with time.

Table 3.2

Number of stocks in each portfolio and sub-period

Portfolio	Period 1	Period 2	Period 3	Period All
FTSE 100	40	70	99	99
FTSE 250	56	102	197	197
FTSE Small Cap	44	93	159	159
FTSE Fledgling	41	79	112	112

Period 1 refers to the test period 01-Jan-1965 to 31-Dec-1979. Period 2 refers to the test period 01-Jan-1980 to 31-Dec-1994. Period 3 refers to the test period 01-Jan-1995 to 30-Jun-2009. Period All refers to the test period 01-Jan-1965 to 30-Jun-2009.

Smaller companies are of specific interest. The literature on moving average trading applied to UK smaller companies is tiny. Belaire-Franch and Opong (2005) apply the variance ratio test of Lo and MacKinlay (1988, 1989) and the non-parametric variance ratio test of Wright (2000) to the value weighted FTSE 100, FTSE 250, FTSE 350 and FTSE All Share indices. The null hypothesis of a random walk is rejected for the FTSE 250 index from January-1986 to September-1997 and for the FTSE All Share index from January-1978 to September-1997. The null hypothesis is not rejected for the FTSE 100 index. Evidence against the FTSE 350 index is mixed. Rejections are due to positive serial dependence. More recently, Hung *et al.* (2009) apply the same tests along with the multiple variance ratio test of Chow and Denning (1993) to the value weighted FTSE 250 and FTSE Small Cap indices from January-1986 to October-2005. The null hypothesis of a random walk is not rejected for the FTSE 250 index but is rejected for the FTSE Small Cap index. As before, rejections are due to positive serial dependence. Failure to reject the null hypothesis for the FTSE 250 index is hard to interpret. An intuitive explanation is that the additional data is so negatively serially dependent that it cancels out the positive serial dependence of the previous data. Why this does not also apply to the FTSE Small Cap index is not clear however. Nevertheless, albeit extremely thin, there

is at least some evidence that the price dynamics of smaller companies are different to those of larger companies and that these dynamics might be suited to moving average trading.

The only study known to apply moving average trading to UK smaller companies is Bokhari *et al.* (2005). The authors test various trading rules using a random sample of stocks drawn from the FTSE 100, FTSE 250 and FTSE Small Cap indices for the period January-1987 to July-2002 and find that predictive ability increases as company size decreases. There is also evidence that the predictive ability of the FTSE Small Cap companies has not declined over time. There is no evidence of profitability however. Simulations in the presence of costs are not profitable. The authors also note that for the FTSE Small Cap companies, the dominant factor stopping the trading rules from being profitable is the size of the bid-ask spread. The trading rules are likely to be profitable if the predictability exhibited by the FTSE Small Cap companies was combined with the bid-ask spread for larger companies. A typical spread for the FTSE Small Cap is quoted as 10.7%. It is not clear if this is representative of the spread used to simulate the trading rules. If it is, it is 3 times larger than in Table 3.1. However, the authors qualify this by noting firstly that the spread is stock dependent and, secondly, that it varies considerably from stock to stock.

3.2 Simulation issues

The moving average used to simulate the trading rules is the exponentially weighted moving average, for which a thorough introduction is Makridakis *et al.* (1998):

$$\begin{aligned}
 ma_0^n(p_0) &= p_0 \\
 ma_t^n(p_t) &= \alpha p_t + (1 - \alpha) ma_{t-1}^n(p_t), \quad \alpha = 2 / (n + 1)
 \end{aligned}
 \tag{3.8}$$

All simulations include the spread. All buys occur at the ask price and all sells occur at the bid price. Buys and sells take place on the same day as the buy/sell signals. A problem with moving average trading is that it is prone to drift if the price does not change for prolonged periods, which sometimes results in false buy/sell signals. To counteract this, all simulations strip out duplicate prices before simulating the trading rules. The buy/sell signals are always refitted to the original price series however. The moving average is always calculated using the close price. Hence, all buy/sell signals are in response to changes in the close price. It is also assumed that there is sufficient liquidity such that each trade executes immediately, pays the spread and has no price impact.

3.3 Test method

A general problem with the trading rules is how to test for statistical significance given their trade distributions. The trade distribution was defined in Chapter 1 as the set of trades T that result from applying a trading rule. The problem is that the theoretical distribution of the test statistics underlying the trade distributions is unknown. This also holds for measures derived as functions of the trade distributions. A popular approach in this situation is to *bootstrap* the test statistics. The bootstrap is due to Efron (1979) and is a computationally intensive method that estimates the distributional properties underlying a sample by re-sampling the empirical or observed distribution to estimate the parameters of interest.

For hypothesis tests, the advantage of the bootstrap is that provided it is possible to specify the distribution under the null hypothesis, statistical inference is straightforward in that the significance of the test statistics can be estimated from the data without having to know the underlying data generating process. All tests for statistical significance are hypothesis tests

and all hypothesis tests are bootstrap tests. The tests are described in detail below. For each test, the number of bootstrap replications B is 500 and the random number generator is *ran2* in Press *et al.* (1996).

Test 1 – Win rate

This is the test used for the win rate in Section 4.1.

The test is described in detail in Efron and Tibshirani (1998) p. 221.

Method

The method is to simulate each stock using the price crossover and trade reduction rules and moving averages in the range $n = 2, 3, \dots, 250$. The trades for each stock for each n are then pooled to give the trade distribution for that n . The trade distribution is the set of trades T that result from applying a trading rule.

Null hypothesis

Let F and G denote the trade distributions for the price crossover and trade reduction rules. The null hypothesis is that the trade distributions for the price crossover and trade reduction rules share the same probability generating function. That is, $H_0 : F = G$.

Test statistic

The test statistic is the difference in the win rates. Under the null hypothesis of no difference, the win rate is the same for the price crossover and trade reduction rules. The win rate is:

$$\frac{n_w}{n_w + n_l} = 1 - \frac{n_l}{n_w + n_l} \quad (3.9)$$

Let \bar{p} and \bar{q} denote the win rates for the price crossover and trade reduction rules. The test statistic is $t(\mathbf{x}) = \bar{q} - \bar{p}$.

Computation of the bootstrap test statistic

Denote the pooled sample as $\mathbf{x} = F + G$ and let f and g denote the number of trades in F and G . The method to compute the bootstrap test statistic is:

1. Draw B samples of size $f + g$ with replacement from \mathbf{x} . Calculate the win rate for the first f samples and for the remaining g samples. Denote the win rate for each sample as \bar{p}^* and \bar{q}^* .
2. Evaluate the test statistic for each sample, $t(\mathbf{x}^{*b}) = \bar{q}^* - \bar{p}^*$, $b = 1, 2, \dots, B$.
3. Calculate the bootstrap p-value as $\#\{t(\mathbf{x}^{*b}) \geq t_{obs}\} / B$ where $t_{obs} = t(\mathbf{x})$, the observed value of the test statistic.

Test 2 – Mean return per trade

This is the test used for the mean return per trade in Section 4.2.

The test is described in detail in Efron and Tibshirani (1998) p. 224.

Method

The method is to simulate each stock using the price crossover and trade reduction rules and moving averages in the range $n = 2, 3, \dots, 250$. The trades for each stock for each n are then pooled to give the trade distribution for that n . The trade distribution is the set of trades T that result from applying a trading rule.

Null hypothesis

Let F and G denote the trade distributions for the price crossover and trade reduction rules. The null hypothesis is that the trade distributions for the price crossover and trade reduction rules have the same mean. That is, $H_0 : \bar{F} = \bar{G}$.

Test statistic

The test statistic is the difference in means. Under the null hypothesis of no difference, the mean return per trade is the same for the price crossover and trade reduction rules. The mean return per trade is defined in Chapter 1. Let \bar{p} and \bar{q} denote the mean return per trade for the price crossover and trade reduction rules. The test statistic is $t(\mathbf{x}) = \bar{q} - \bar{p}$.

Computation of the bootstrap test statistic

Efron and Tibshirani (1998) stress the importance of sampling under the null hypothesis. To test the null hypothesis it is necessary for F and G to have the same mean. Let f and g denote the number of trades in F and G . The method to compute the bootstrap test statistic is:

1. Let $\tilde{f}_i = f_i - \bar{f} + \bar{z}$, $i = 1, 2, \dots, f$ and let $\tilde{g}_i = g_i - \bar{g} + \bar{z}$, $i = 1, 2, \dots, g$ where \bar{f} and \bar{g} are the means of F and G and \bar{z} is the mean of the combined sample.
2. Draw B samples of size f with replacement from $\tilde{f}_1, \tilde{f}_2, \dots, \tilde{f}_f$ and B samples of size g with replacement from $\tilde{g}_1, \tilde{g}_2, \dots, \tilde{g}_g$. Denote the mean of each sample as \bar{p}^* and \bar{q}^* .
3. Evaluate the test statistic for each sample, $t(\mathbf{x}^{*b}) = \bar{q}^* - \bar{p}^*$, $b = 1, 2, \dots, B$.
4. Calculate the bootstrap p-value as $\#\{t(\mathbf{x}^{*b}) \geq t_{obs}\} / B$ where $t_{obs} = t(\mathbf{x})$, the observed value of the test statistic.

3.4 Summary

This chapter has introduced the data and described the bootstrap hypothesis tests used to test the trade reduction rule in the next chapter. The main concern has been to derive an estimate of the spread as a measure of cost. The spread is important when simulating the trading rules

since failure to account for costs overstates the return and can lead to erroneous conclusions regarding profitability. More specifically, failure to account for costs overstates the win rate and the mean return per trade, both of which are tested next.

Chapter 4

Trade Reduction Rule

This chapter introduces the trade reduction rule. The trade reduction rule is based on the idea of allowing a trade to run and extends the price crossover rule to increase the mean return per trade by keeping as many winning trades as possible at the same time as reducing the number of losing trades. Results for buying and selling are different. The mean return per trade for the long only trade reduction rule is consistently higher than for the price crossover rule. This is true for all portfolios for all test periods. The average increase in the mean return per trade is 84%. The mean return per trade for the short only trade reduction rule is consistently higher than for the price crossover rule for the FTSE Small Cap and FTSE Fledgling portfolios only. Evidence against the null hypothesis of no difference in the means of the trades generated by the trade reduction and price crossover rules is mixed. It is not possible to conclusively reject the null hypothesis for all portfolios for all test periods. However, there are large numbers of trading rules that reject the null hypothesis and where failure to reject the null hypothesis is otherwise marginal. Overall, there is sufficient reason to conclude that the trade reduction rule behaves as expected even if it does not follow that the difference in the mean return per trade is consistently statistically significant. The remainder of the chapter is organised as follows. Section 4.1 derives the trade reduction rule. Section 4.2 presents the test results for the null hypothesis of no difference in the means of the trades generated by the trade reduction and price crossover rules. Section 4.3 offers conclusions.

4.1 Trade reduction rule

This section defines the trade reduction rule. The trade reduction rule has two parts. The first part is the bounded moving average. The bounded moving average maps the moving average onto the range $[-1, 1]$. This has a number of applications including testing the distribution of buy/sell signals for deviation from a null model, testing the position of the buy/sell signals for differences in profitability as well as allowing for trading rules that respond to patterns in the path of the moving average. This latter application underpins the trade reduction rule and the second part describes how knowing the order and position of minima and maxima can be used to generalise the price crossover rule in such a way as to transform the mean return per trade. Section 4.1.1 defines the bounded moving average. Section 4.1.2 describes the trade reduction algorithm.

4.1.1 Bounded moving average

The intuition behind bounding the moving average is that the more the price falls, the more the moving average falls and the more the price and the moving average are *low* relative to before. Similarly, the more the price rises, the more the moving average rises and the more the price and the moving average are *high* relative to before. Mapping the moving average onto $[-1, 1]$ where -1 corresponds to low and 1 corresponds to high constrains the moving average to move within a fixed range. This makes it possible to measure the position of the moving average as defined by its location within $[-1, 1]$. The bounding algorithm is:

1. Compute the n -day moving average of the price $ma_t^n(p_t)$
2. Set the minimum of the price and $ma_t^n(p_t)$ to $b_t^- = \min(p_t, ma_t^n(p_t))$
3. Set the maximum of the price and $ma_t^n(p_t)$ to $b_t^+ = \max(p_t, ma_t^n(p_t))$

Since $b_t^- \leq p_t \leq b_t^+$ for all t , it holds that $ma_t^n(b_t^-) \leq ma_t^n(p_t) \leq ma_t^n(b_t^+)$:

4. Compute the n -day moving average of b_t^- and set this to the lower bound $ma_t^n(b_t^-)$
5. Compute the n -day moving average of b_t^+ and set this to the upper bound $ma_t^n(b_t^+)$

The moving average is then normalised to fit $[-1, 1]$:

6. Define the bounded moving average as $\phi(ma_t^n(p_t)) = \frac{2(ma_t^n(p_t) - ma_t^n(b_t^-))}{ma_t^n(b_t^+) - ma_t^n(b_t^-)} - 1$

Figure 4.1 plots an example. The minima and maxima of Figure 4.1 are especially important. Mathematically, when the price crosses its average, the moving average changes direction. If the price cuts up through its average from below, the moving average changes direction from falling to rising. Similarly, if the price cuts down through its average from above, the moving average changes direction from rising to falling. This means that the minima and maxima in the moving average of Figure 4.1 are identical to the buy/sell signals generated by the price crossover rule defined in Chapter 1. Further, the minima and maxima in the moving average are identical to the minima and maxima in the bounded moving average. Hence, the position of the buy/sell signals is known. This information can be used to remodel the price crossover rule by changing the way it responds to the position of the buy/sell signals and is the approach

underlying the trade reduction rule. Given that their buy/sell signals are not generated by the price crossing its average, there is no equivalent for the other types of trading rule defined in Chapter 1. For the other types of trading rule, the bounded moving average could also follow a path similar to Figure 4.1 as the result of a single trade. This means that it is hard to identify patterns that can be defined as general case. However, it is possible to capture something of the nature of the other types of trading rule provided not all minima and maxima result in a buy/sell signal. A feature of the trade reduction rule is that it achieves this without a second moving average.

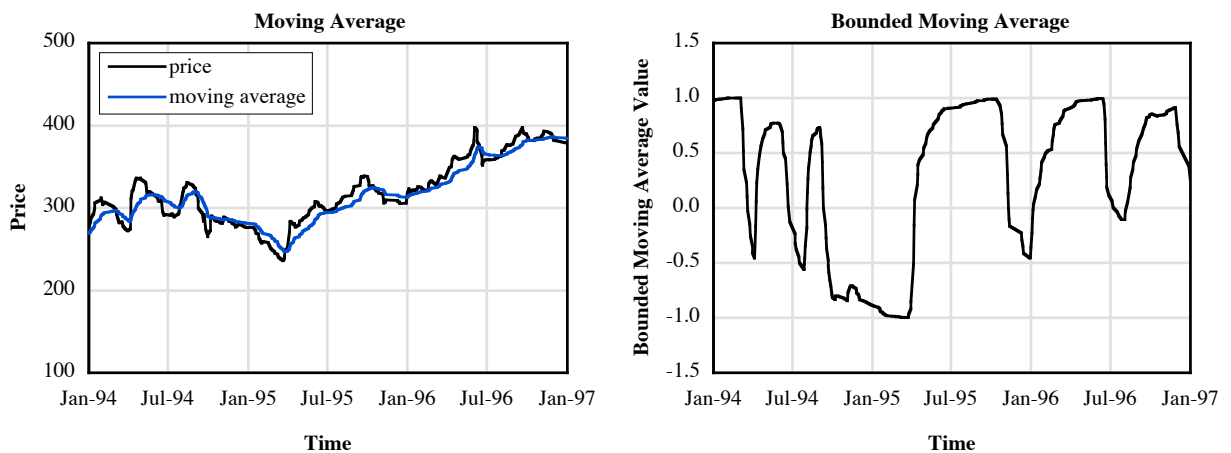


Figure 4.1 Example of the moving average and its mapping to the bounded moving average.

Figure 4.2 plots an example frequency distribution of the position of the minima and maxima within $[-1, 1]$. The distribution is non-normal and heavily weighted in the left and right tails. Minima tend to occur close to -1 and maxima tend to occur close to 1 . The reason for this is that the bounded moving average is defined self-referentially. The bounded moving average measures the extent to which the price goes up and down in terms of the similarity between

the moving average and itself. As a result, the clustering in the tails is due to persistence in direction. The longer the price continues in the same direction, the more likely it is for the moving average to resemble itself. The more likely it is for the moving average to resemble itself, the more likely it is for the bounded moving average to approach and settle on ± 1 . The U-shape is due to the bounding algorithm. The same U-shape is found in a random walk with the difference in dynamics reflected in differences in frequency. Random walks are found to be less dense in the left and right tails and more heavily weighted through the middle.

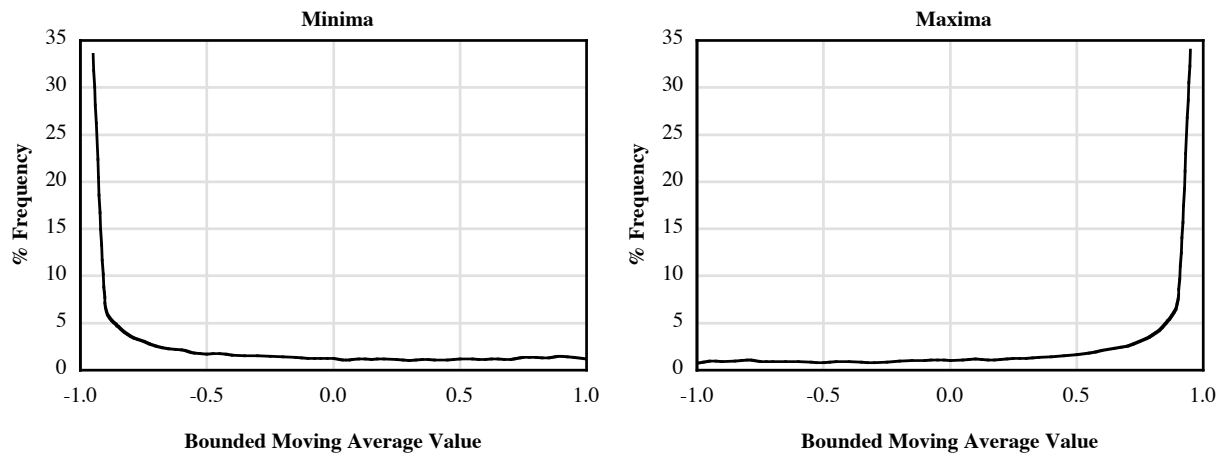


Figure 4.2 Example frequency distribution of the position of minima and maxima.

4.1.2 Trade reduction algorithm

To explain the trade reduction rule, it helps to first define the price crossover rule. As stated, the minima and maxima in the moving average are identical to the buy/sell signals generated by the price crossover rule. Define a local minimum $lmin_t$ and a local maximum $lmax_t$ as:

$$lmin_t = ma_{t-2}^n(p_t) \geq ma_{t-1}^n(p_t) < ma_t^n(p_t) \quad (4.1)$$

$$lmax_t = ma_{t-2}^n(p_t) \leq ma_{t-1}^n(p_t) > ma_t^n(p_t) \quad (4.2)$$

The long only price crossover rule can then be written as:

$$q_1 = 0$$

///// RULE

$$q_{t+1} = q_t$$

$$\text{IF } (lmin_t) \text{ THEN } q_{t+1} = 1$$

$$\text{IF } (lmax_t) \text{ THEN } q_{t+1} = 0 \quad (4.3)$$

Similarly, the short only price crossover rule can be written as:

$$q_1 = 0$$

///// RULE

$$q_{t+1} = q_t$$

$$\text{IF } (lmin_t) \text{ THEN } q_{t+1} = 0$$

$$\text{IF } (lmax_t) \text{ THEN } q_{t+1} = -1 \quad (4.4)$$

The nature of the price crossover rule is that a trade always comprises consecutive minima and maxima. The trade reduction rule generalises the price crossover rule to allow a trade to comprise multiple minima and maxima. This is achieved by exploiting the ordering and the position of the minima and maxima. The objective is to transform the mean return per trade by keeping as many winning trades as possible at the same time as reducing the number of

losing trades. The mean return per trade that results from applying the price crossover rule was defined in Chapter 1 and is repeated here for convenience:

$$\bar{R} = \frac{(n_w \times \bar{W}) + (n_l \times \bar{L})}{n_w + n_l} \quad (4.5)$$

More generally, denote the mean return per trade by k/n . The approach is to transform the mean return per trade such that:

$$\frac{k}{m} > \frac{k}{n} \quad (4.6)$$

For positive k , (4.6) follows whenever $m < n$. Let R_t denote the return for the current trade.

The long only trade reduction rule is:

$$\begin{aligned} q_1 &= 0 \\ \text{previousLMin}_0 &= 0 \\ \text{///// RULE} \\ q_{t+1} &= q_t \\ \text{previousLMin}_t &= \text{previousLMin}_{t-1} \\ \text{IF } (lmin_t \text{ AND } q_{t+1} = 0) \text{ THEN } q_{t+1} &= 1 \\ \text{IF } (lmax_t \text{ AND } R_t > 0) \text{ THEN } q_{t+1} &= 0 \\ \text{IF } (lmin_t) \text{ THEN } \text{previousLMin}_t &= \phi(ma_t^n(p_t)) \\ \text{IF } (\phi(ma_t^n(p_t)) \leq \text{previousLMin}_t) \text{ THEN } q_{t+1} &= 0 \end{aligned} \quad (4.7)$$

In the event of a local minimum, the trade reduction rule opens a new long trade if it has not already done so and saves the position of the local minimum for when it is needed later. The next event is a local maximum. When a local maximum occurs, the trade is either profitable or it is not. If it is profitable, it is closed. This keeps as many winning trades as possible. The trading process then restarts from the next local minimum. The sequence of events and their relationship with the mean return per trade is:

$$lmin_{t-i}, lmax_t \rightarrow n_w = n_w + 1, \sum_{x \in W} x = \sum_{x \in W} x + R_t \quad (4.8)$$

If the trade is not profitable, rather than take a loss, the trade reduction rule waits for the next buy/sell signal event. The reason this is said to be a buy/sell signal event is that normally the next buy/sell signal would be a local minimum. The trade reduction rule also generates sell signals whenever the bounded moving average cuts down through the position of the previous local minimum. Hence, the source of the next buy/sell signal is not known until it occurs. If the bounded moving average cuts down through the position of the previous local minimum before the next local minimum occurs, the trade is closed and is expected to result in a loss. If the trade was not profitable at the price of the previous local maximum, it is unlikely to be profitable after the moving average has fallen from this. The trading process then restarts from the next local minimum. The sequence of events is as shown below where a struck out signal indicates that it is not acted on:

$$lmin_{t-j}, \cancel{lmax_{t-i}}, \phi(ma_t^n(p_t)) \leq \text{previous}LMin_t \rightarrow n_l = n_l + 1, \sum_{x \in L} x = \sum_{x \in L} x + R_t \quad (4.9)$$

The reasons for closing the trade when the bounded moving average cuts down through the position of the previous local minimum are twofold. Firstly, because the buy/sell signals of the price crossover rule are ordered as minima and maxima, if a local minimum does not occur before the bounded moving average cuts down through the position of the previous local minimum, it is known that the next local minimum will occur at a position less than this. Since the objective is to keep as many winning trades as possible, closing the trade at this point ensures that the trading process restarts from the next local minimum. This means that not too many buy signals are ignored and that the next trade will be kept should it turn out to be a winning trade. Secondly, there is an open long trade when the moving average is falling. In theory, the moving average could fall forever resulting in bankruptcy. This is undesirable. More generally, if the moving average continues falling, this could result in significant losses. Closing the trade when the bounded moving average cuts down through the position of the previous local minimum is a simple method by which to control the magnitude of the losses. However, the best that can be achieved is that the loss does not get any bigger than it already is.

The remaining situation is where a local minimum occurs before the bounded moving average cuts down through the position of the previous local minimum. The trade reduction rule does nothing in this case other than to update the position of the previous local minimum to that of the current local minimum. The trading process then continues from the next local maximum. Because there has been a local minimum followed by a local maximum followed by a local minimum but where the latter two signals have not been acted on, the effect is to reduce the number of losing trades relative to the price crossover rule by one. This is also reflected in the sum of the losing trades. The sequence of events and their relationship with the mean return per trade relative to the price crossover rule is:

$$lmin_{t-j}, \cancel{lmass_{t-i}}, \cancel{lmin_t} \rightarrow n_l = n_l - 1, \sum_{x \in L} x = \sum_{x \in L} x - R_l \quad (4.10)$$

There are several elements at work here. Firstly, updating the position of the previous local minimum to the current local minimum ensures that all trades are closed relative to the most recent local minimum. This has some similarities to a trailing stop loss where the stop loss is adjusted to be the most recent local minimum. If the trade ends up as a losing trade, updating the position of the local minimum ensures that the trade will be exited earlier, which helps to control the magnitude of the losses. Secondly, the trade reduction rule is defined recursively. Each time the pattern defined by (4.10) occurs, it is succeeded by either (4.8), (4.9) or (4.10). If succeeding events are unbiased in that (4.8) sometimes occurs, some trades will end up as winning trades. Not only does the trade reduction rule reduce the number of losing trades, it also increases the number of winning trades. It is not possible to create winning trades from nothing however. Consequently, the increase in the number of winning trades is likely to be small. Lastly, the position of the previous local minimum is always less than the position of the current local minimum. Because both minimums are associated with the prices at which they occur, there is also the possibility that the price level of the previous local minimum is less than the price level of the current local minimum. If so, there is the possibility that the difference in price level will bias the outcome in favour of a winning trade. Figure 4.3 plots some example trades (with annotation) for the long only trade reduction rule in terms of the path of the bounded moving average.



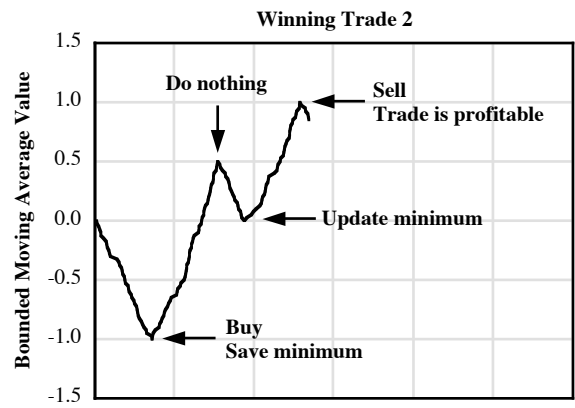



Figure 4.3 Example trades for the long only trade reduction rule in terms of the path of the bounded moving average.

The short only trade reduction rule is the converse:

$$q_1 = 0$$

$$previousLMax_0 = 0$$

///// RULE

$$q_{t+1} = q_t$$

$$previousLMax_t = previousLMax_{t-1}$$

$$\text{IF } (lmax_t \text{ AND } q_{t+1} = 0) \text{ THEN } q_{t+1} = -1$$

$$\text{IF } (lmin_t \text{ AND } R_t > 0) \text{ THEN } q_{t+1} = 0$$

$$\text{IF } (lmax_t) \text{ THEN } previousLMax_t = \phi(ma_{t-1}^n(p_t))$$

$$\text{IF } (\phi(ma_t^n(p_t)) \geq previousLMax_t) \text{ THEN } q_{t+1} = 0 \quad (4.11)$$

The question then is what to expect from the trade reduction rule. Firstly, the trade reduction rule should result in fewer trades than the price crossover rule. Secondly, there should also be more winning trades and less losing trades. The trade reduction rule should have a higher *win rate* where the win rate is the ratio of winning trades to the total number of trades. Tests find that this is in fact the case. On average, the trade reduction rule generates 7% more winning trades and 46% fewer losing trades than the price crossover rule. Bootstrap tests of the null hypothesis of no difference in the win rate are soundly rejected. This is true for all portfolios for all test periods long and short. Figure 4.4 plots a typical example of the test results for the number of winning and losing trades using the FTSE 100 portfolio for the period 01-January-1965 to 30-June-2009. Similar results are found for all portfolios for all test periods. The trade reduction rule significantly transforms the trade distribution of the price crossover rule due, primarily, to the decrease in the number of losing trades. While the increase in the number of

winning trades is also a contributory factor, the decrease in the number of losing trades is the dominant element.

This is an encouraging result and suggests that the trade reduction rule should have an effect on the mean return per trade. However, it can be seen from Figure 4.4 that the trade reduction rule has little impact on the distribution of winning trades. This is to be expected. Given that the winning trades are due to trending behaviour in the underlying price dynamics, the trade reduction rule cannot be expected to uncover new trends not already uncovered by the price crossover rule. On the other hand, the trade reduction rule does have a significant impact on the distribution of losing trades. The problem is that to facilitate this, the losses incurred by the trade reduction rule are different to the losses incurred by the price crossover rule. For the price crossover rule, losses are always incurred at the most recent local maximum and their magnitude is always determined by the difference in the price level relative to the previous local minimum. For the trade reduction rule, losses are allowed to include multiple minima and maxima and are incurred after the moving average has fallen from the most recent local maximum. It is not clear if this difference is significant. Analysis finds that the difference is not significant. For the long only trade reduction rule, the sum of the losing trades is always less than for the price crossover rule. This is true for all portfolios for all test periods. For the short only trade reduction rule, the sum of the losing trades is sometimes greater than for the price crossover rule although this is true only for the FTSE 100 and FTSE 250 portfolios for test period 2. The difference is also negligible. Figure 4.5 plots a typical example of the test results for the sum of the winning and losing trades using the FTSE 100 portfolio for the period 01-January-1965 to 30-June-2009. Similar results are found for all portfolios for all test periods. The trade reduction rule does not generate significantly higher losses than the price crossover rule and nor does it generate significantly lower profits. The trade reduction

rule can therefore be expected to have a similar return to the price crossover rule. Assuming the return is positive, given the overall decrease in the number of trades, the trade reduction rule can also be expected to have a higher mean return per trade. A higher mean return per trade also implies higher breakeven costs, which may or may not be significant in terms of outperforming buy and hold.

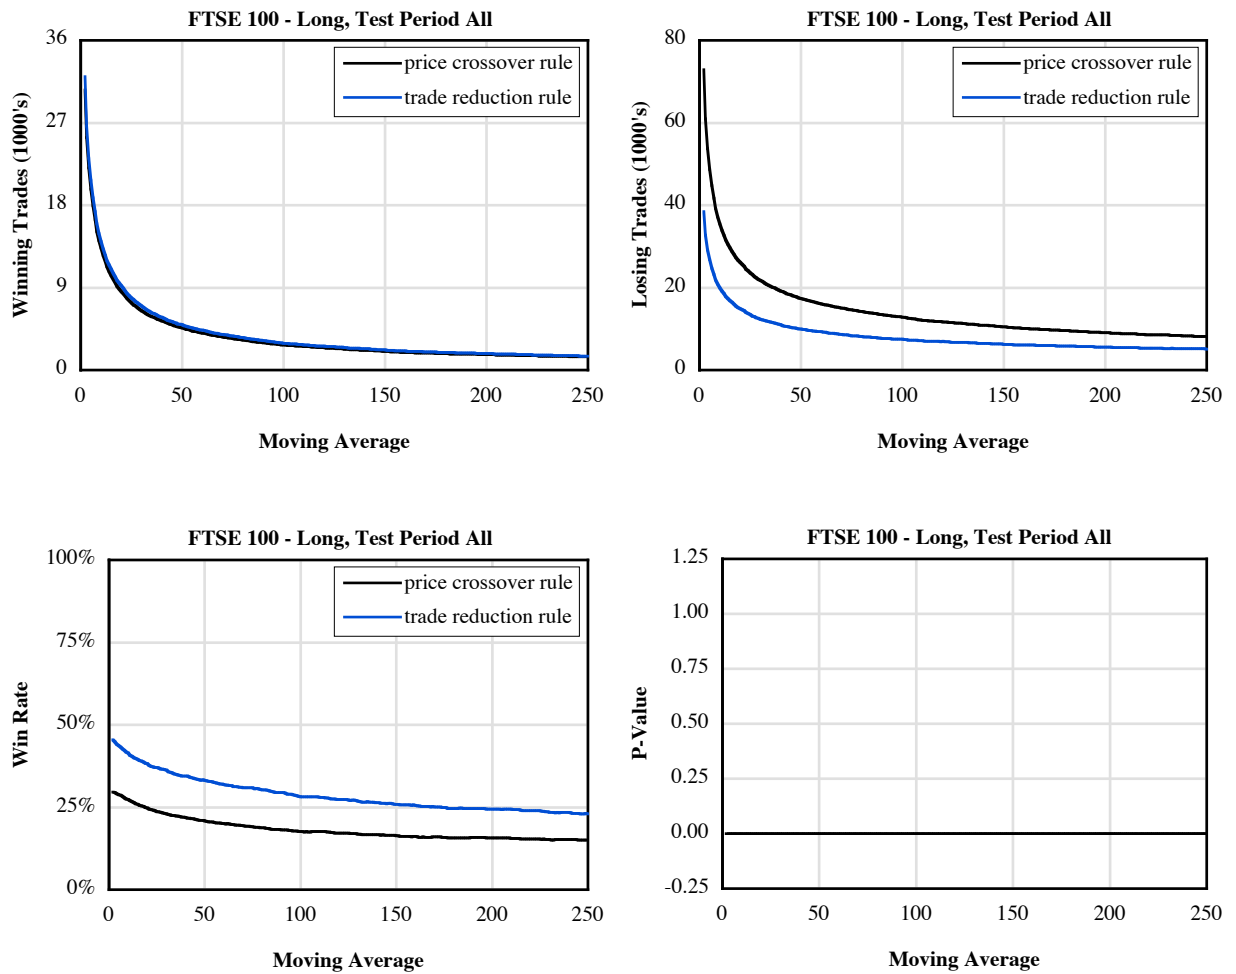


Figure 4.4a Number of winning and losing trades for the long only price crossover and trade reduction rules for the FTSE 100 portfolio for the period 01-January-1965 to 30-June-2009. The p-values are the bootstrap results for the null hypothesis of no difference in the win rate.

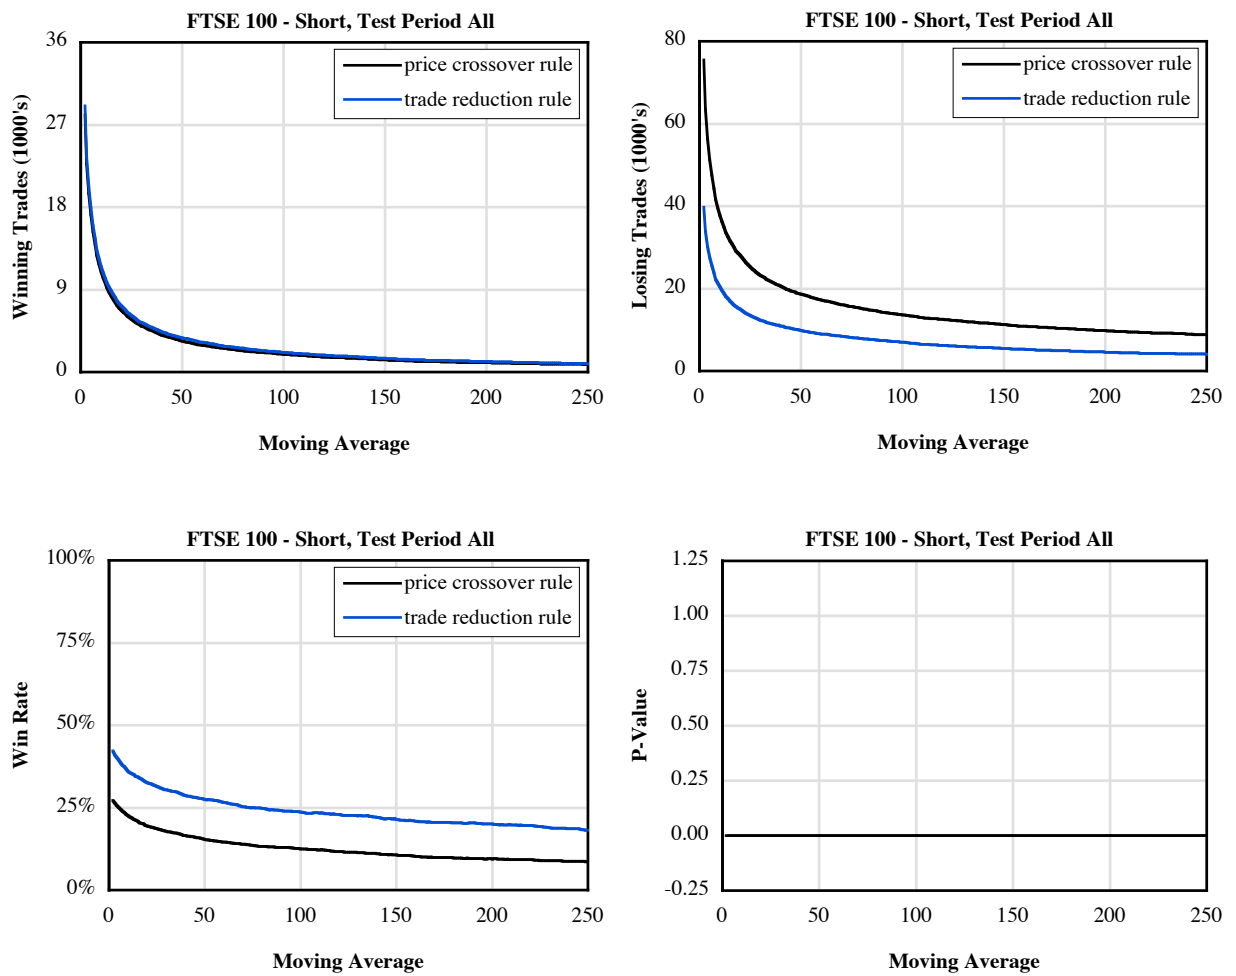


Figure 4.4b Number of winning and losing trades for the short only price crossover and trade reduction rules for the FTSE 100 portfolio for the period 01-January-1965 to 30-June-2009.

The p-values are the bootstrap results for the null hypothesis of no difference in the win rate.

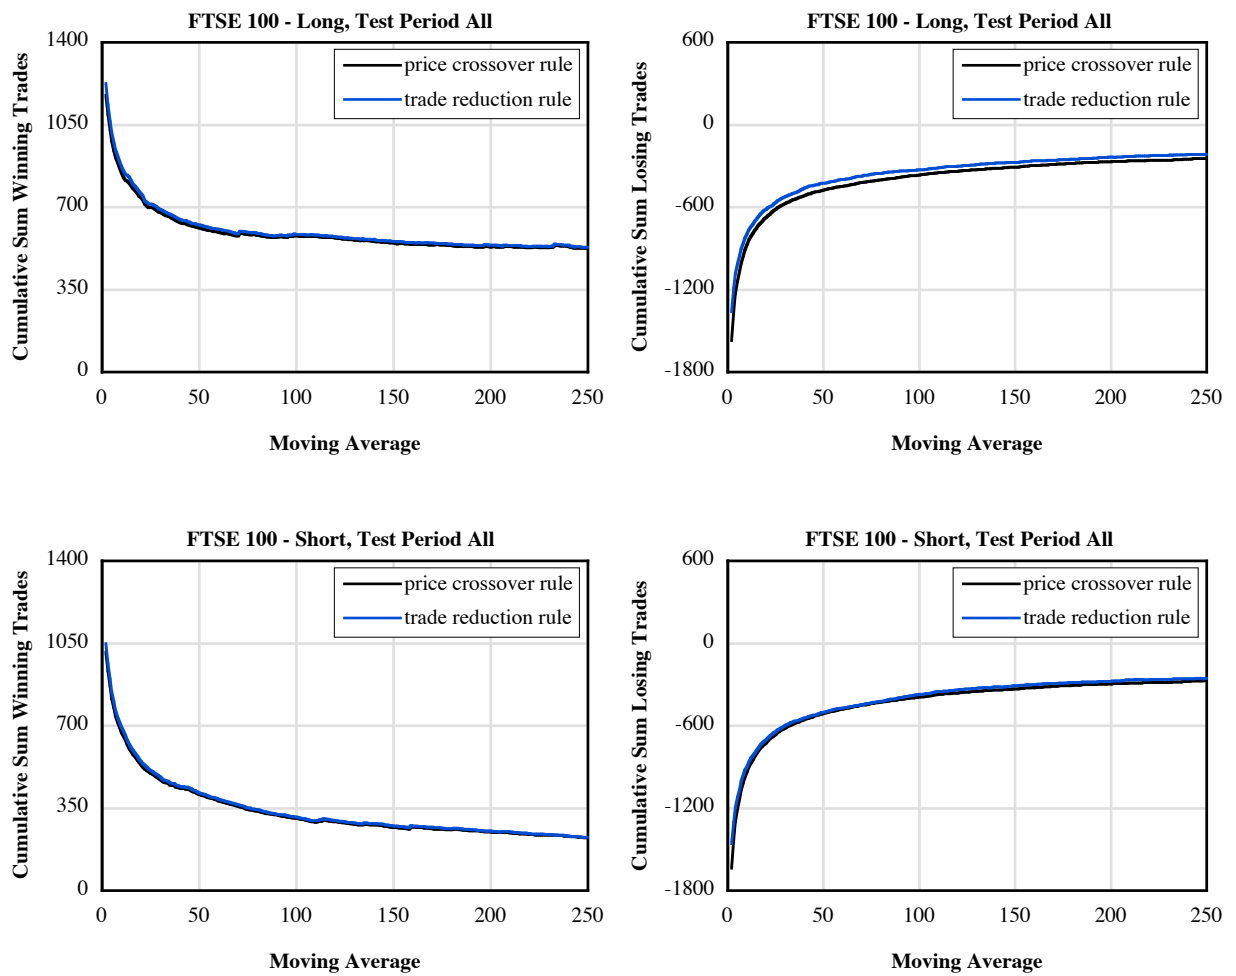


Figure 4.5 Sum of the winning and losing trades for the long only and short only price crossover and trade reduction rules for FTSE 100 portfolio for the period 01-January-1965 to 30-June-2009.

4.2 Results for the mean return per trade

This section presents the first set of simulation results for the trade reduction rule. The aim is to test the null hypothesis of no difference in the means of the trades generated by the price crossover and trade reduction rules. The method is to simulate each stock individually using moving averages in the range $n = 2, 3, \dots, 250$ and to combine the trades for each stock for each n to obtain the set of trades for that n . The trade distributions are then bootstrapped to test the null hypothesis. The test is described in Chapter 3.¹ Note that the test is indicative of whether the trade distributions have the same mean only. It does not test the actual return to the trading rules. Results for the return are presented later. Section 4.2.1 presents the results for the long only price crossover and trade reduction rules. The mean return per trade for the trade reduction rule is consistently higher than for the price crossover rule. This is true for all portfolios for all test periods. There is also evidence that the mean return per trade is lowest post-1990 and that the mean return per trade increases as company size decreases. However, evidence against the null hypothesis is mixed. In general, there are large numbers of trading rules that reject the null hypothesis and where failure to reject the null hypothesis is otherwise marginal. Overall, there is sufficient reason to conclude that the trade reduction rule behaves as expected even though it does not follow that the difference in the mean return per trade is consistently statistically significant. Section 4.2.2 presents the results for the short only price crossover and trade reduction rules. Results for the short only trading rules are more variable and there is evidence that the returns to selling are different. The mean return per trade for the trade reduction rule is consistently higher than for the price crossover rule for the FTSE Small Cap and FTSE Fledgling portfolios only. The mean return per trade for the FTSE Small Cap and FTSE Fledgling portfolios is also unusually high. This may be indicative of survivorship

¹ See Section 3.3 (page 46).

bias. As before, evidence against the null hypothesis is mixed. It is also weaker. Nevertheless, in general, the trade reduction rule works much the same for selling as it does for buying with the differences due to the difference in the behaviour of the underlying price dynamics.

4.2.1 Long

Figures 4.6 to 4.9 plot the mean return per trade for the long only price crossover and trade reduction rules.² Bootstrap p-values for the null hypothesis of no difference in the means are also shown. The mean return per trade for the trade reduction rule is consistently higher than for the price crossover rule. This is true for all portfolios for all test periods. The mean return per trade is also lowest in test period 3. This is consistent with the literature where the returns to moving average trading have declined post-1990. The mean return per trade also increases as company size decreases and suggests that prices are more likely to trend as company size decreases.

With reference to the p-values, evidence against the null hypothesis of no difference in the means is mixed. In general, it is not possible to conclusively reject the null hypothesis for all portfolios for all test periods. However, there are large numbers of trading rules that reject the null hypothesis and where failure to reject the null hypothesis is otherwise marginal. This is particularly evident when testing all of the data. The p-values in this case are strong evidence against the null hypothesis although their variability across sub-periods would suggest that a note of caution is perhaps more appropriate. Even so, as a whole, there is sufficient reason to

² A long trade refers to buying in the expectation that prices will rise. The stock is then sold at a higher price than the price for which it was bought. The profit is the sell price minus the buy price. See Section 1.2 (page 4).

conclude that the trade reduction rule behaves as expected even if it does not follow that the difference in the mean return per trade is consistently statistically significant. This suggests that to consistently reject the null hypothesis, the trade reduction rule needs to transform the distributions of both the winning and losing trades together and not just the losing trades on their own.

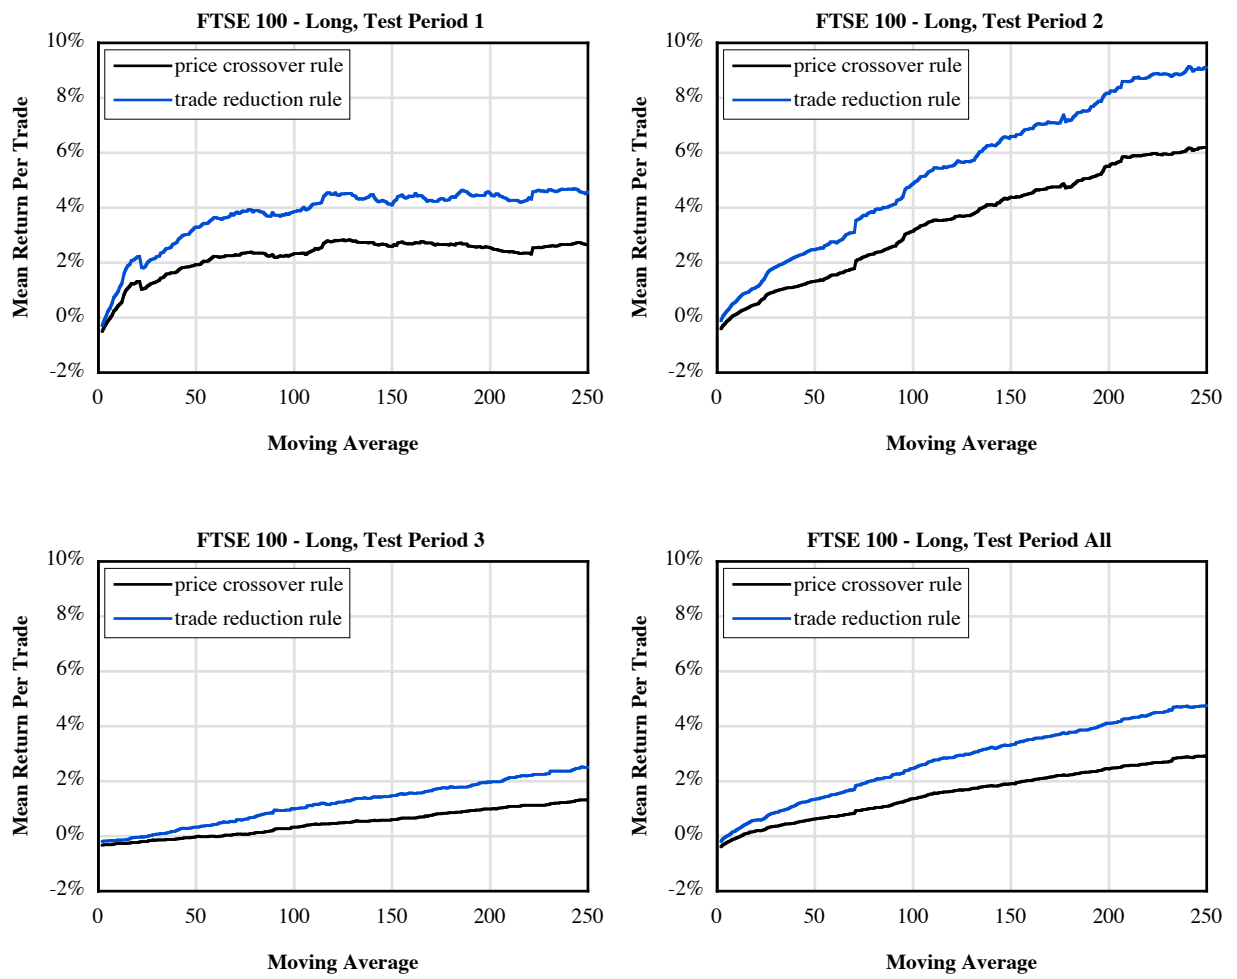


Figure 4.6a Mean return per trade for the long only price crossover and trade reduction rules for the FTSE 100 portfolio.

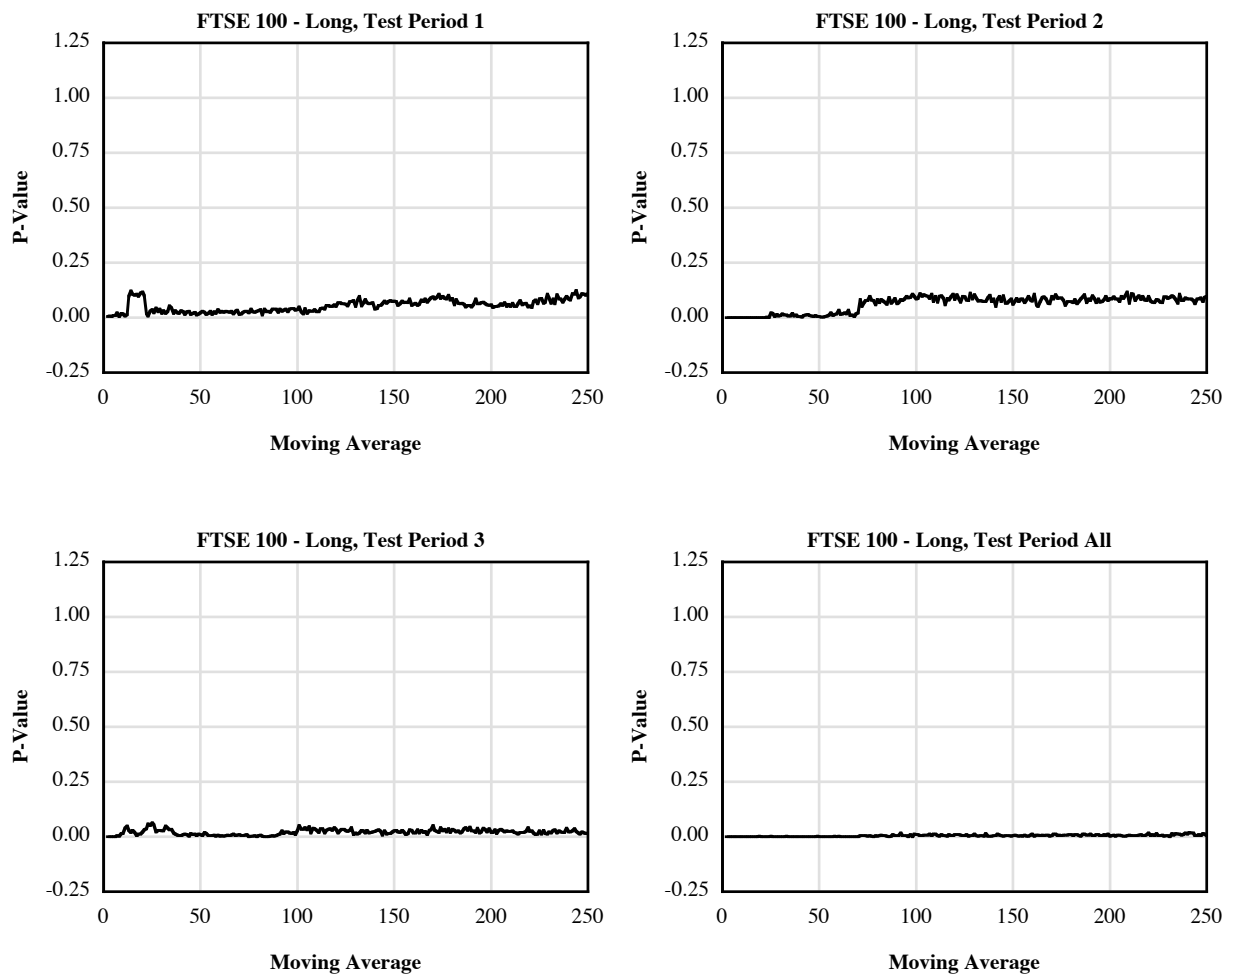


Figure 4.6b Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the long only price crossover and trade reduction rules for the FTSE 100 portfolio.

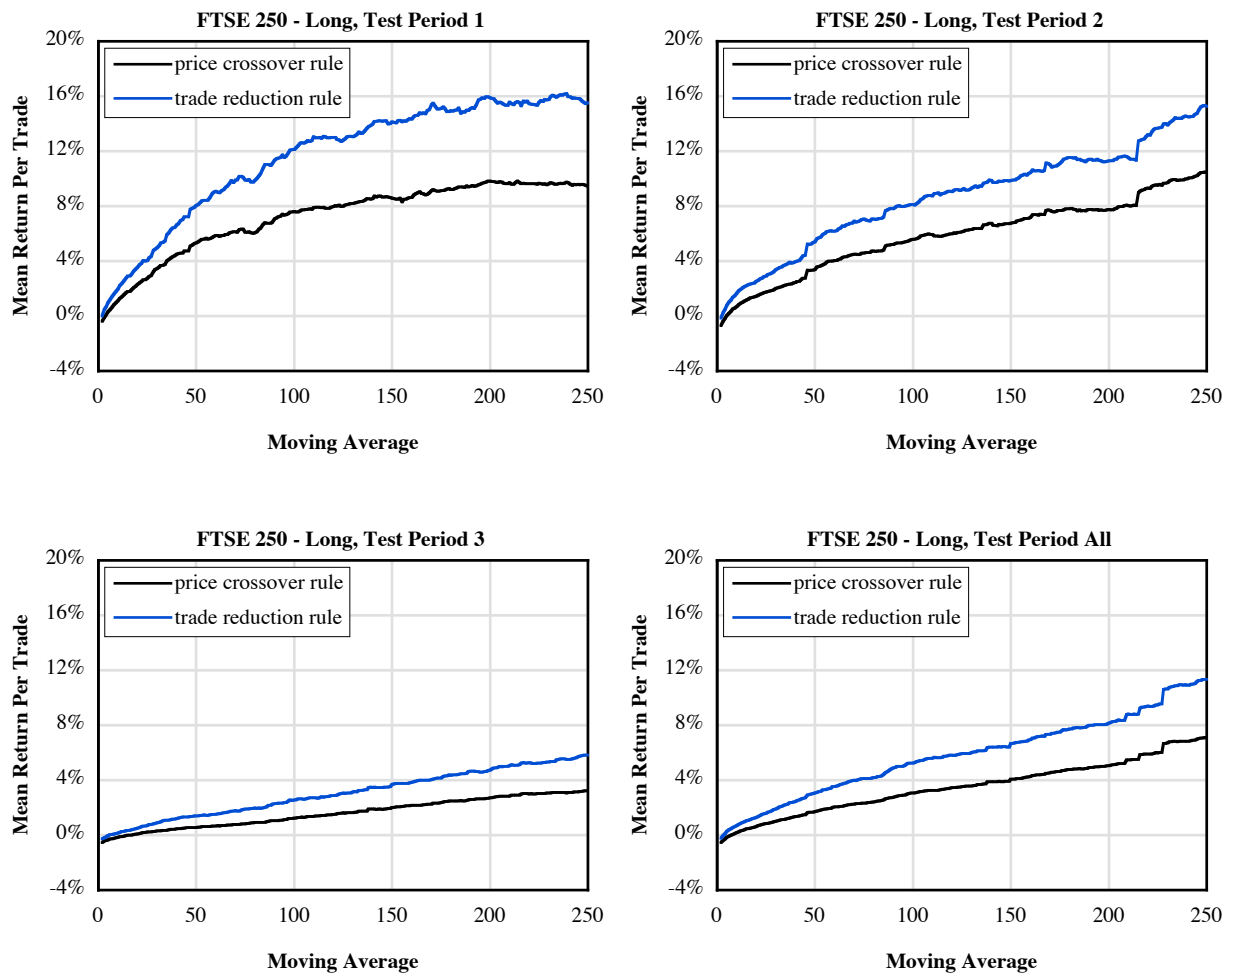


Figure 4.7a Mean return per trade for the long only price crossover and trade reduction rules for the FTSE 250 portfolio.

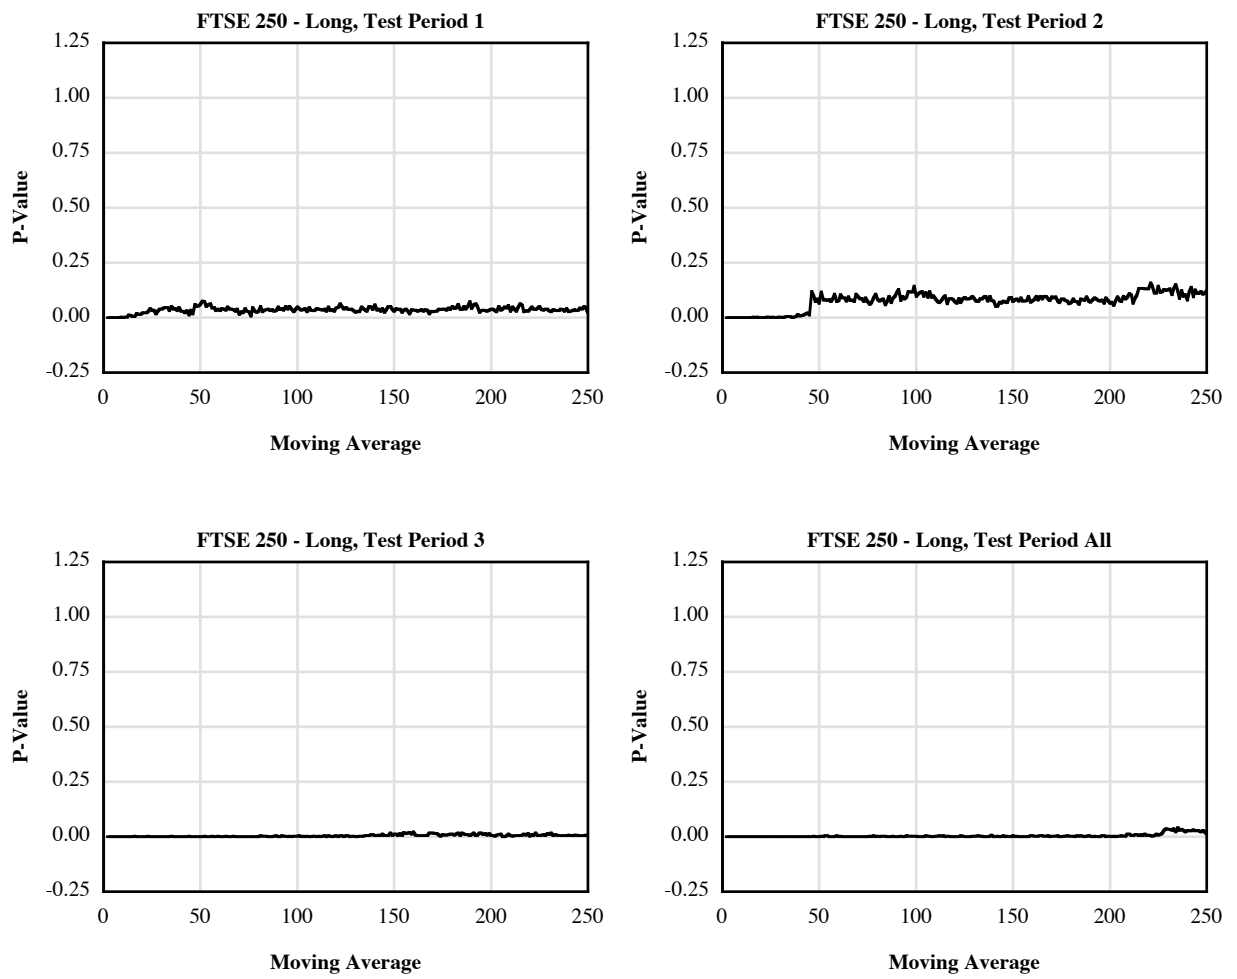


Figure 4.7b Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the long only price crossover and trade reduction rules for the FTSE 250 portfolio.

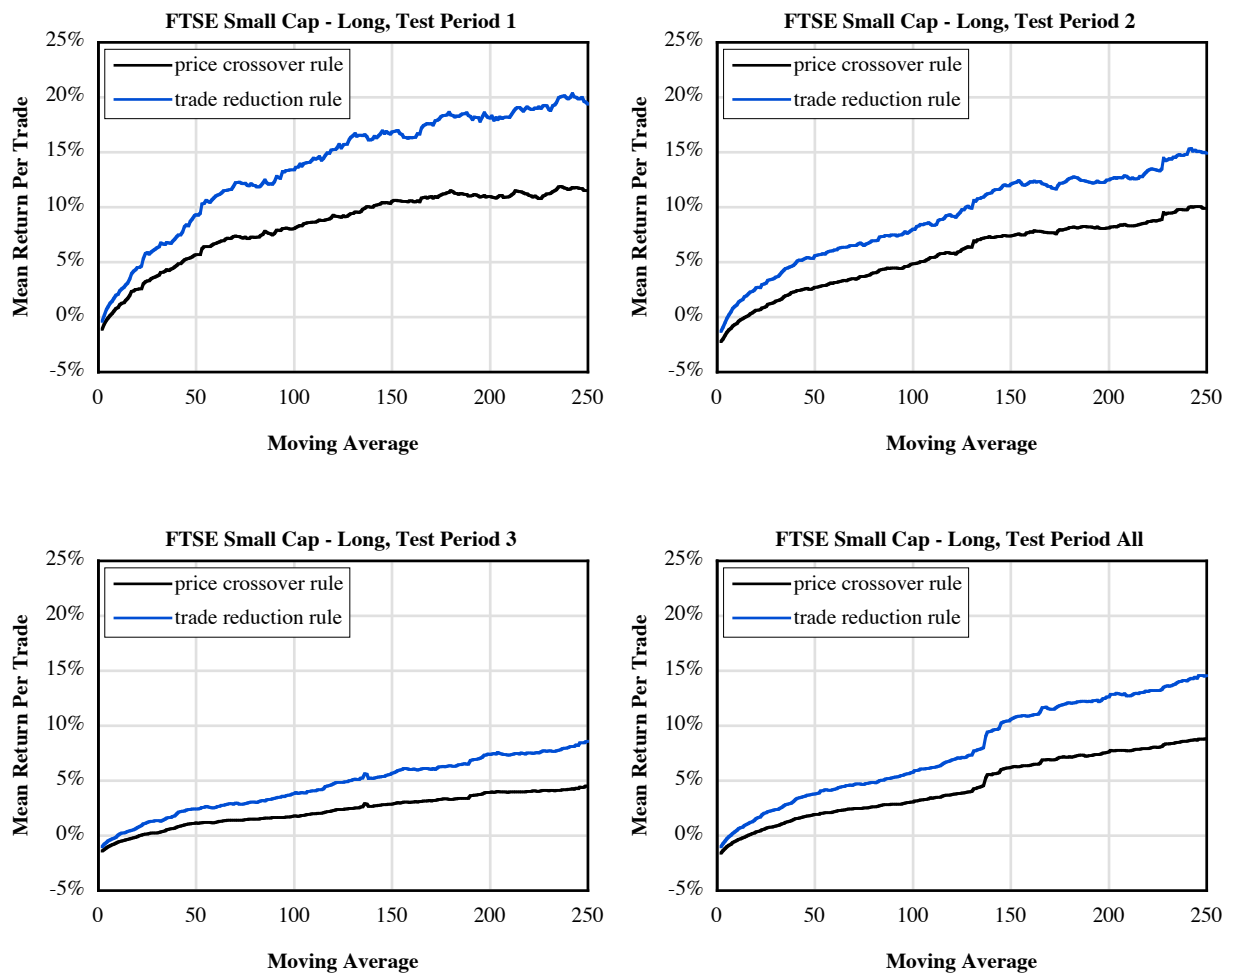


Figure 4.8a Mean return per trade for the long only price crossover and trade reduction rules for the FTSE Small Cap portfolio.

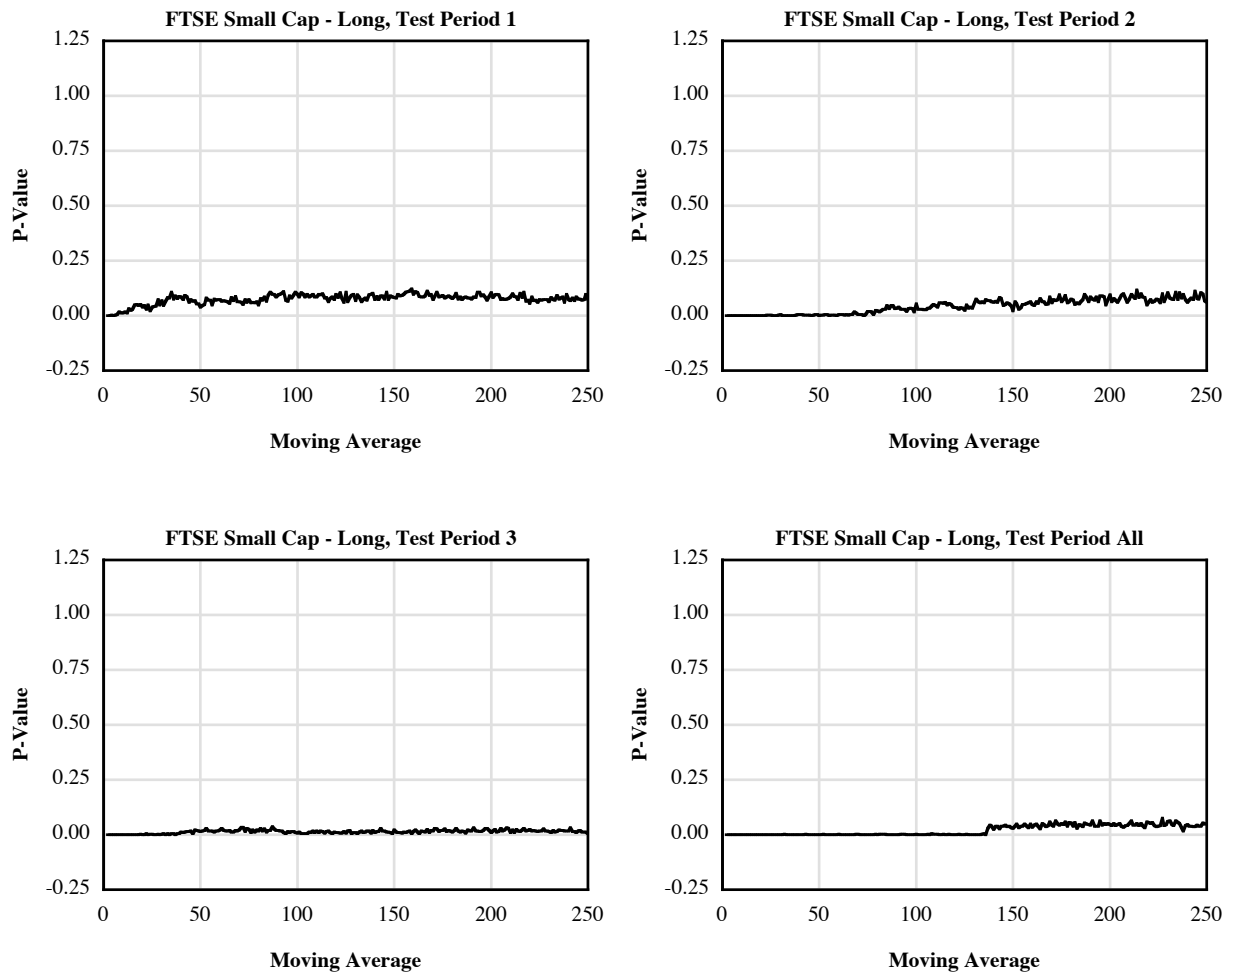


Figure 4.8b Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the long only price crossover and trade reduction rules for the FTSE Small Cap portfolio.

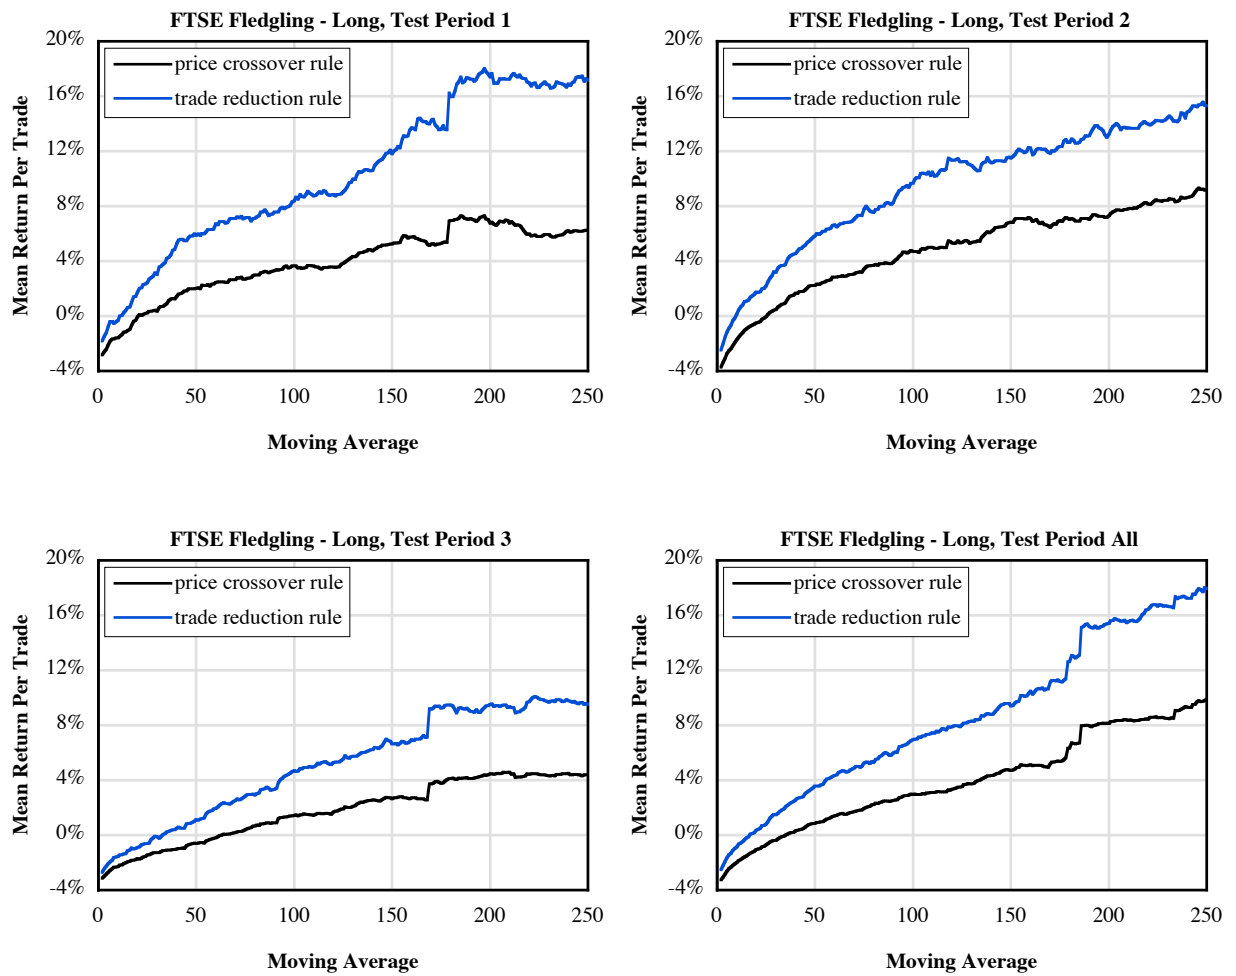


Figure 4.9a Mean return per trade for the long only price crossover and trade reduction rules for the FTSE Fledgling portfolio.

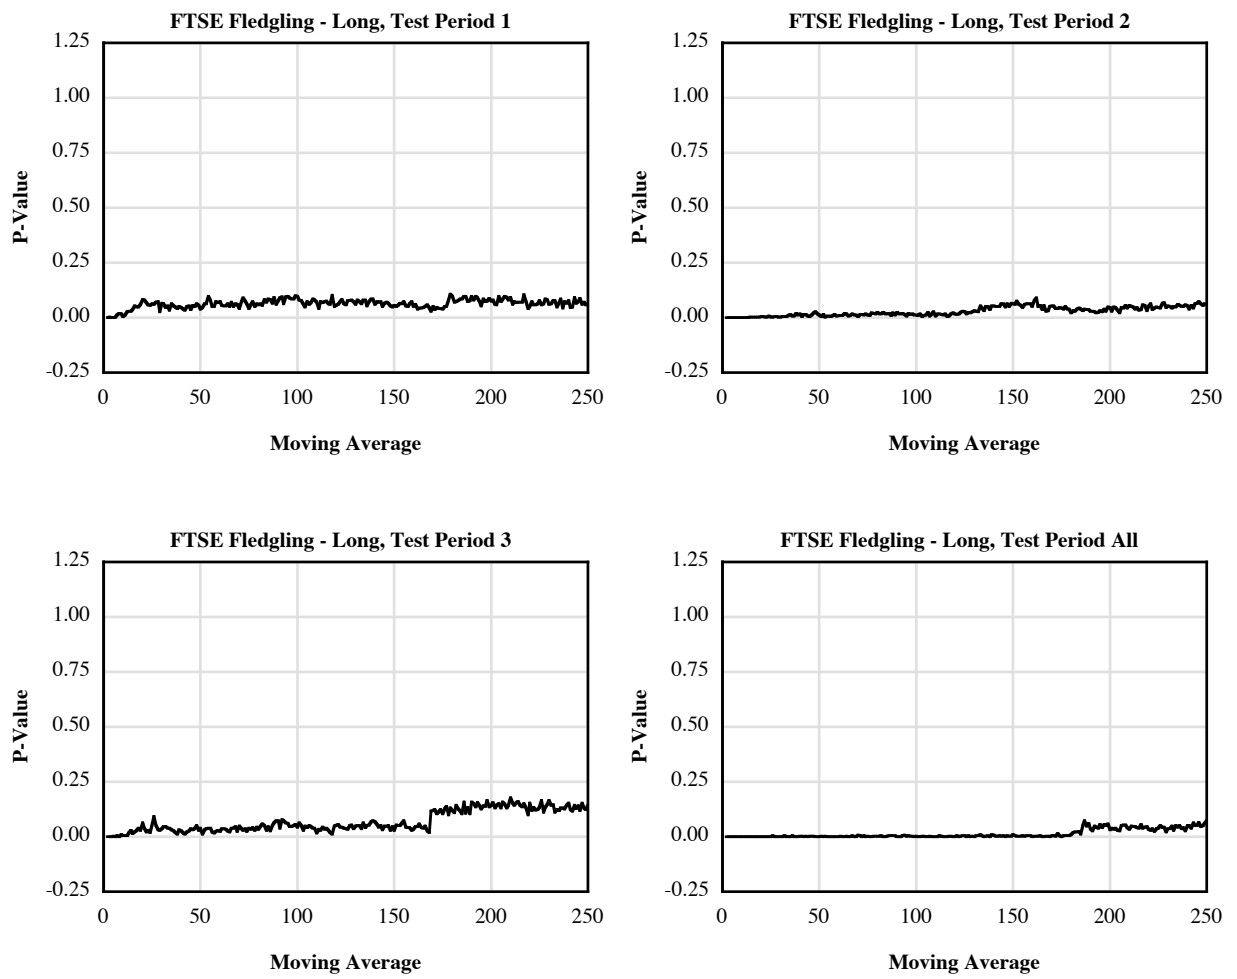


Figure 4.9b Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the long only price crossover and trade reduction rules for the FTSE Fledgling portfolio.

4.2.2 Short

Figures 4.10 to 4.13 plot the mean return per trade for the short only price crossover and trade reduction rules.³ Bootstrap p-values for the null hypothesis of no difference in the means are also shown. The returns to selling are clearly different from the returns to buying. The mean return per trade for the trade reduction rule is consistently higher than for the price crossover rule for the FTSE Small Cap and FTSE Fledgling portfolios only. With the exception of test period 1, there is little evidence that shorting the FTSE 100 portfolio is likely to be profitable. The same observation applies to the FTSE 250 portfolio with the exception of test periods 1 and 3. On the other hand, the opposite is likely to be true for the FTSE Small Cap and FTSE Fledgling portfolios, most notably for test periods 2 and 3. The mean return per trade for the FTSE Small Cap and FTSE Fledgling portfolios during these periods is unusually high. This is a difficult result to interpret. Because the mean return per trade increases as company size decreases, as with buying, it might be that prices for smaller companies are more likely to trend. If so, this appears to be true for prices on the way down as well as on the way up. The FTSE index history of the portfolio constituents is not known though. Since the FTSE Small Cap and FTSE Fledgling portfolios comprise the smallest 2% – 3% of companies that make up the main market but where the portfolios were not constructed until August 2009, looking back in time, the portfolios might be biased by companies whose businesses are in decline but which continue to maintain a listing and are still trading. In other words, the FTSE Small Cap and FTSE Fledgling portfolios are likely to include a number of stocks that are present in the portfolios for the reason that their value has fallen to the point where there is nowhere else for

³ A short trade refers to selling in the expectation that prices will fall. The stock is then bought back at a lower price than the price for which it was sold. The profit is the sell price minus the buy price. See Section 1.2 (page 4).

them to go. If so, the high mean return per trade is at least partly explained by survivorship bias. This tends to be supported by the results for test period 1 where there are fewer stocks in the portfolios and where the general shape of the mean return per trade is consistent with that of the FTSE 100 and FTSE 250 portfolios. A further complication is that some stocks trade as penny shares whereby a small change in the price can generate significant returns. Results for shorting the FTSE Small Cap and FTSE Fledgling portfolios need to be treated with caution particularly with respect to conclusions based on company size alone and to the existence of a negative risk premium. Then again, if it were possible to acquire or at least approximate the FTSE Small Cap or FTSE Fledgling portfolios in real-time, assuming similar results, such a cautionary note would perhaps not apply.

As with buying, evidence against the null hypothesis of no difference in the means is mixed. In general, it is not possible to conclusively reject the null hypothesis for all portfolios for all test periods. Evidence against the null hypothesis is also weaker than for buying. However, there are large numbers of trading rules that reject null hypothesis and where failure to reject the null hypothesis is otherwise marginal. This is most apparent for test period 3. This is an acceptable if not decisive result and it is reasonable to conclude that the trade reduction rule works much the same for selling as it does for buying with the differences attributable to the differences in the behaviour of the underlying price dynamics.

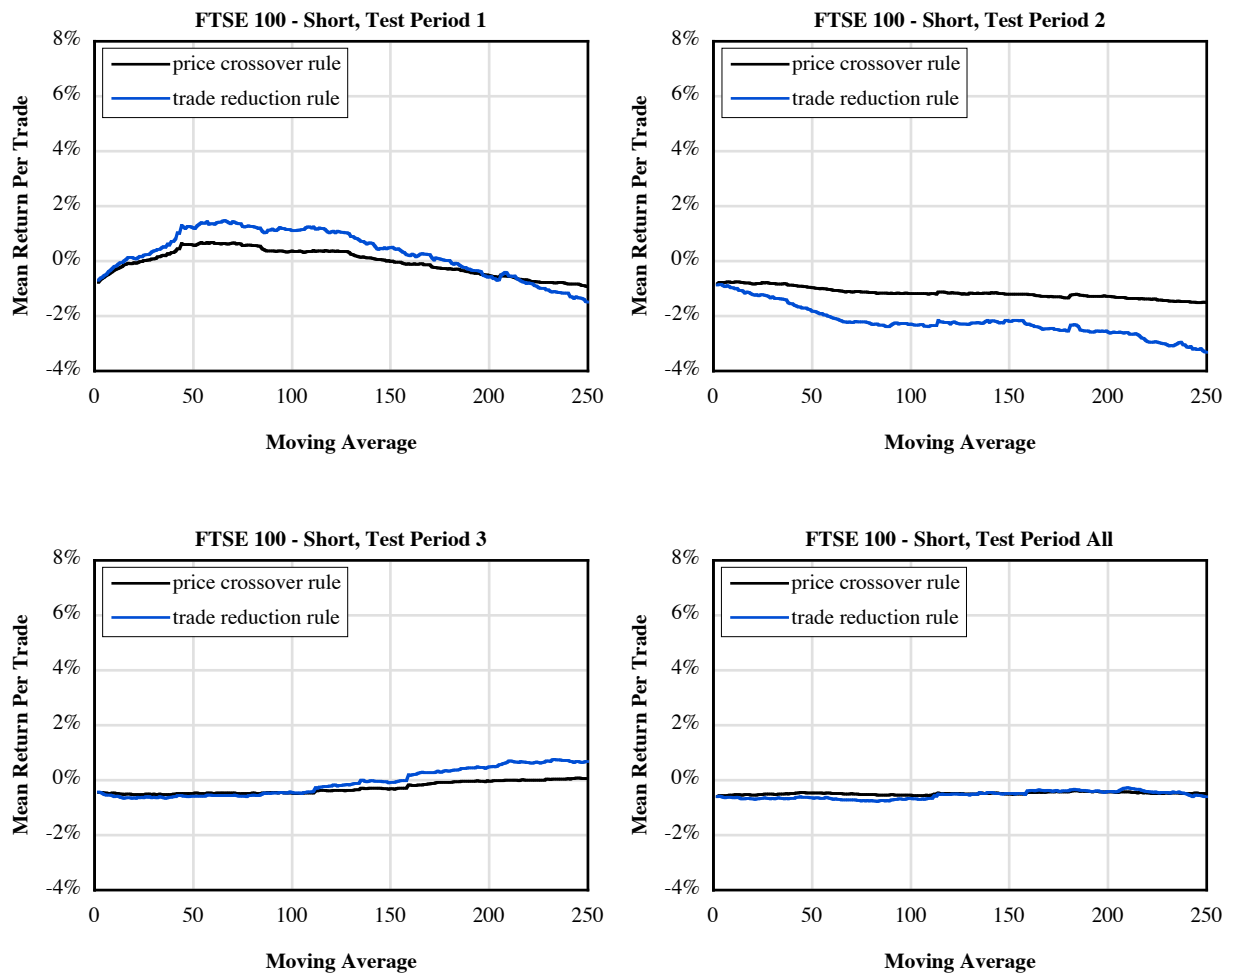


Figure 4.10a Mean return per trade for the short only price crossover and trade reduction rules for the FTSE 100 portfolio.




Figure 4.10b Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the short only price crossover and trade reduction rules for the FTSE 100 portfolio.

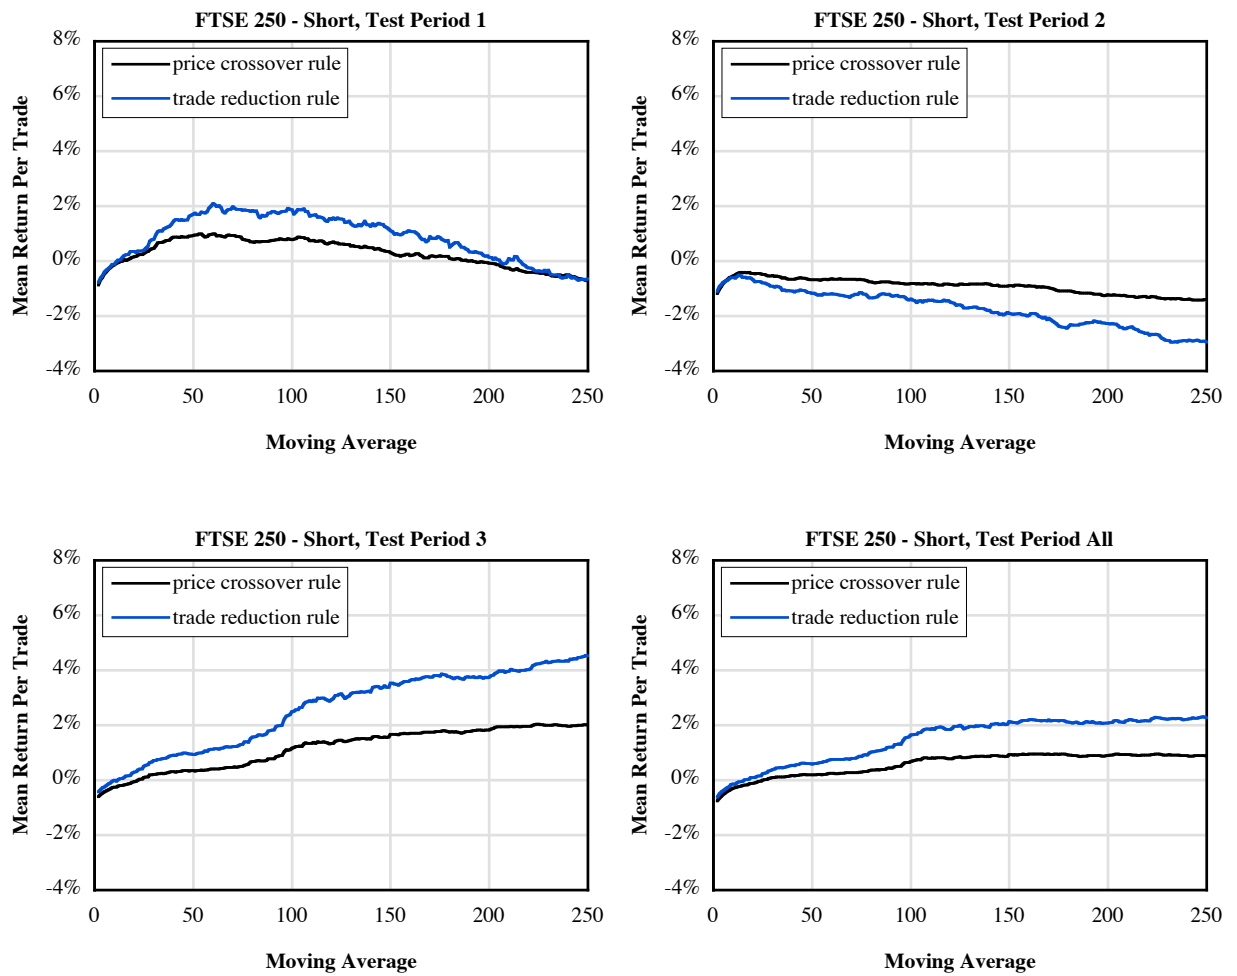


Figure 4.11a Mean return per trade for the short only price crossover and trade reduction rules for the FTSE 250 portfolio.

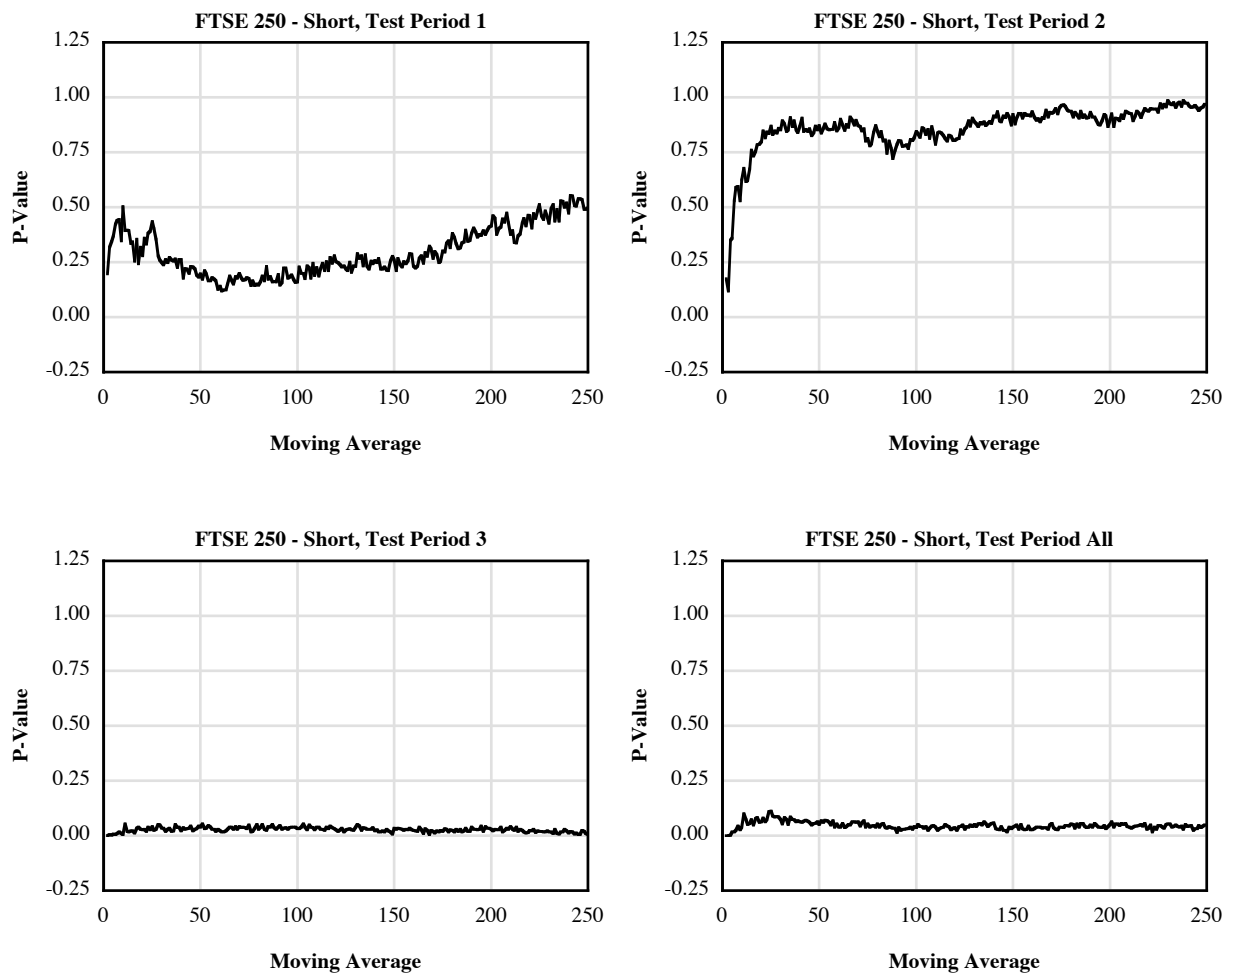


Figure 4.11b Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the long short price crossover and trade reduction rules for the FTSE 250 portfolio.

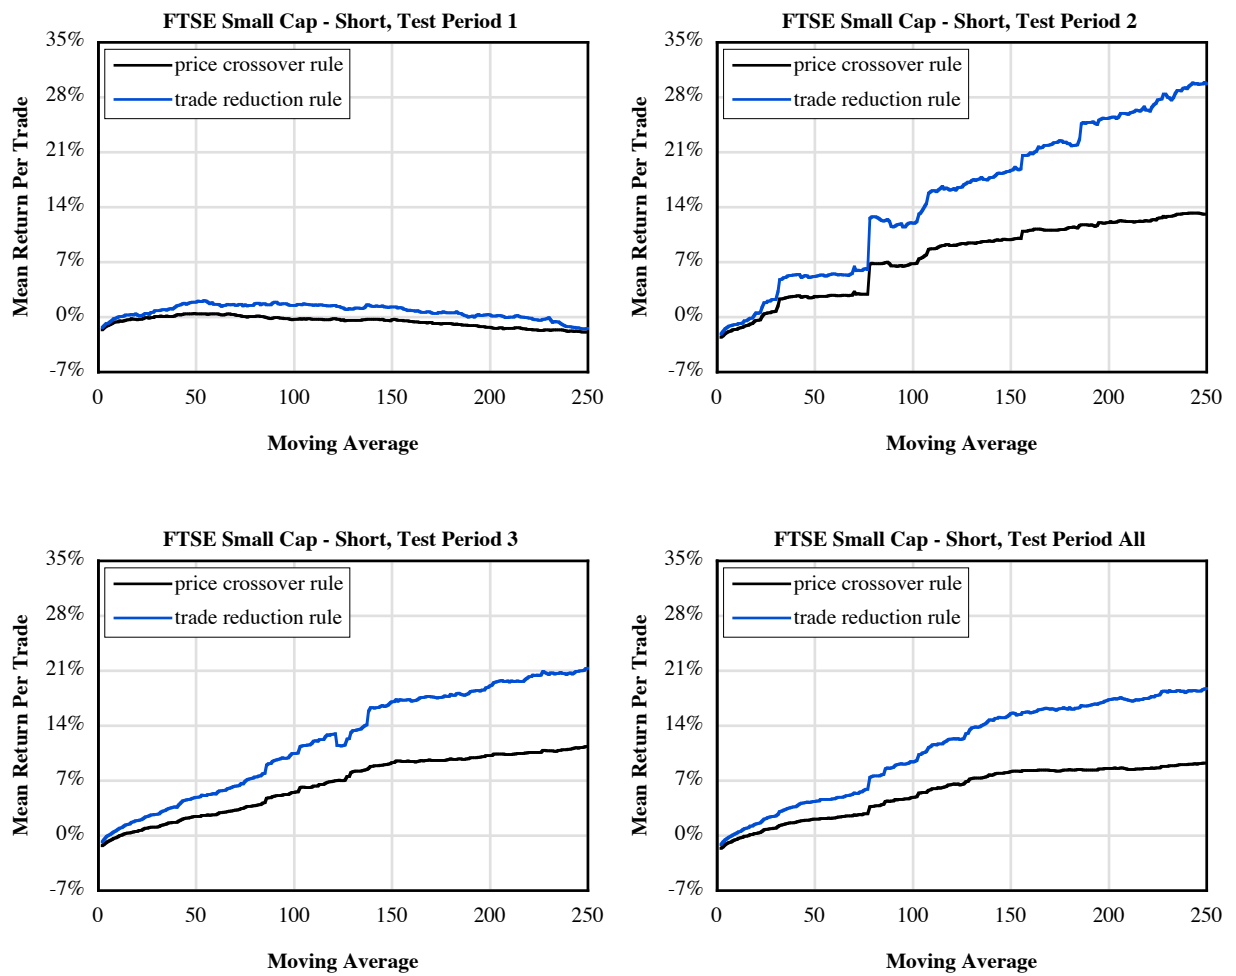


Figure 4.12a Mean return per trade for the short only price crossover and trade reduction rules for the FTSE Small Cap portfolio.

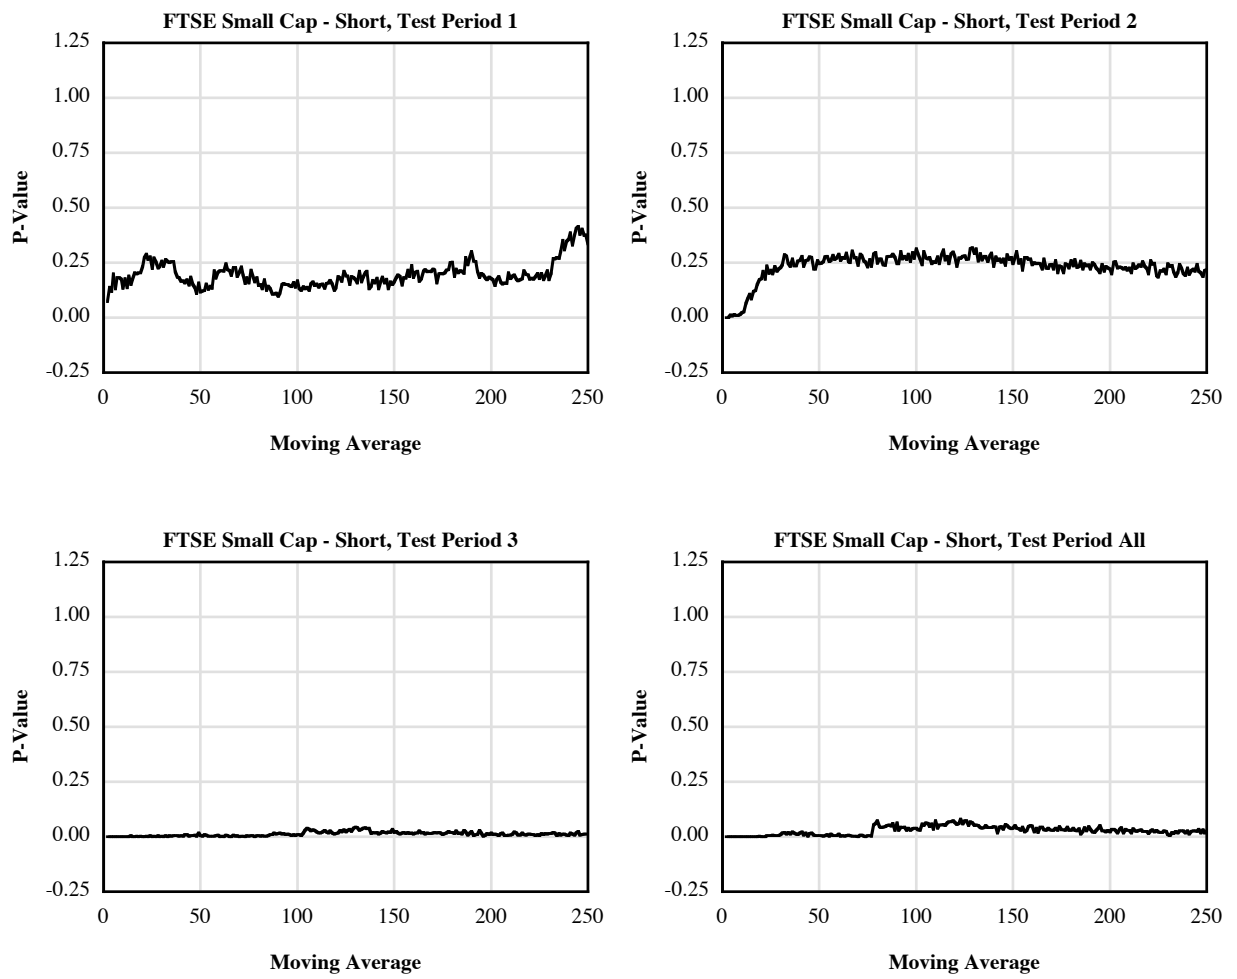


Figure 4.12b Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the short only price crossover and trade reduction rules for the FTSE Small Cap portfolio.

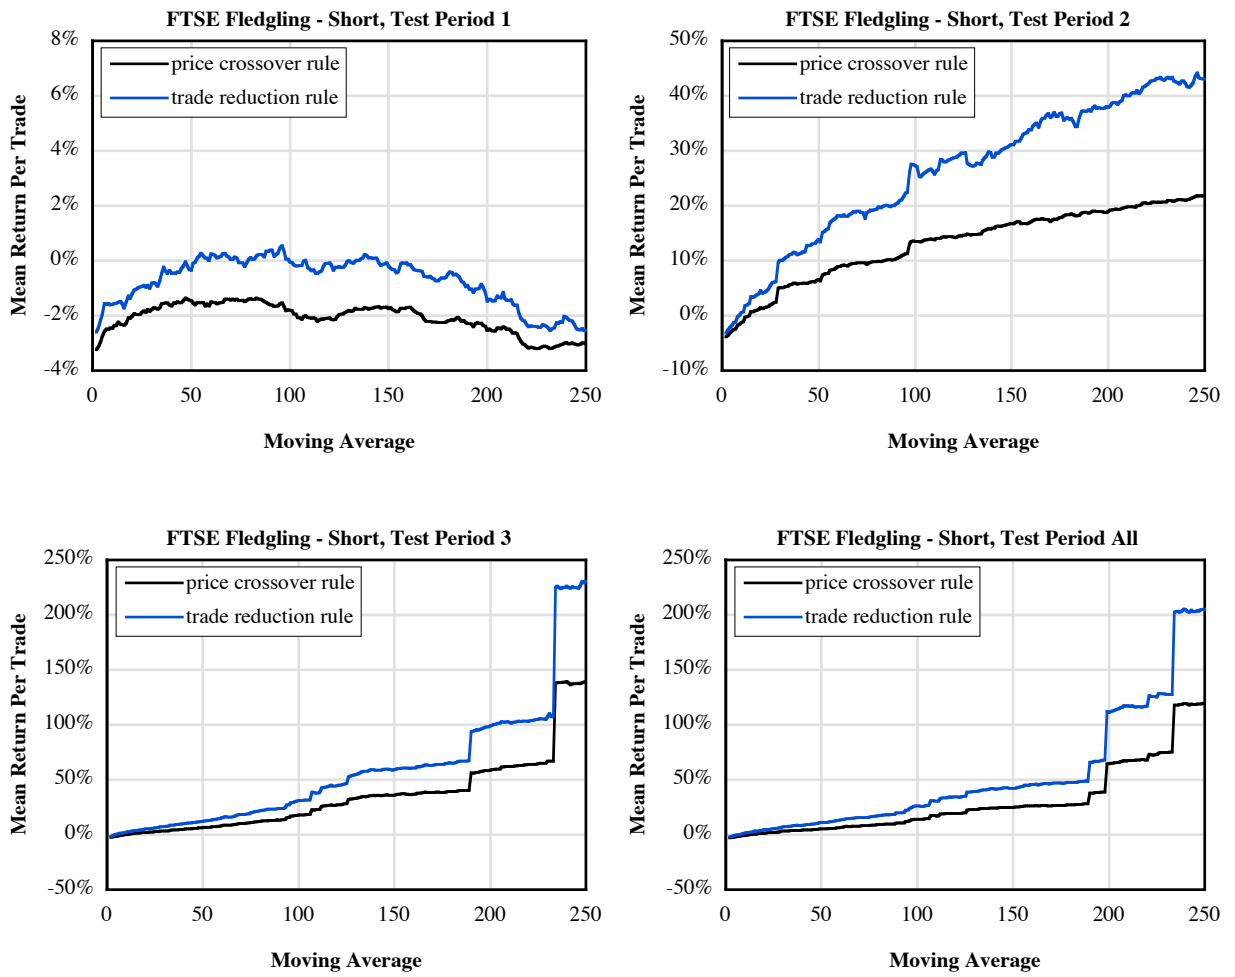


Figure 4.13a Mean return per trade for the short only price crossover and trade reduction rules for the FTSE Fledgling portfolio.

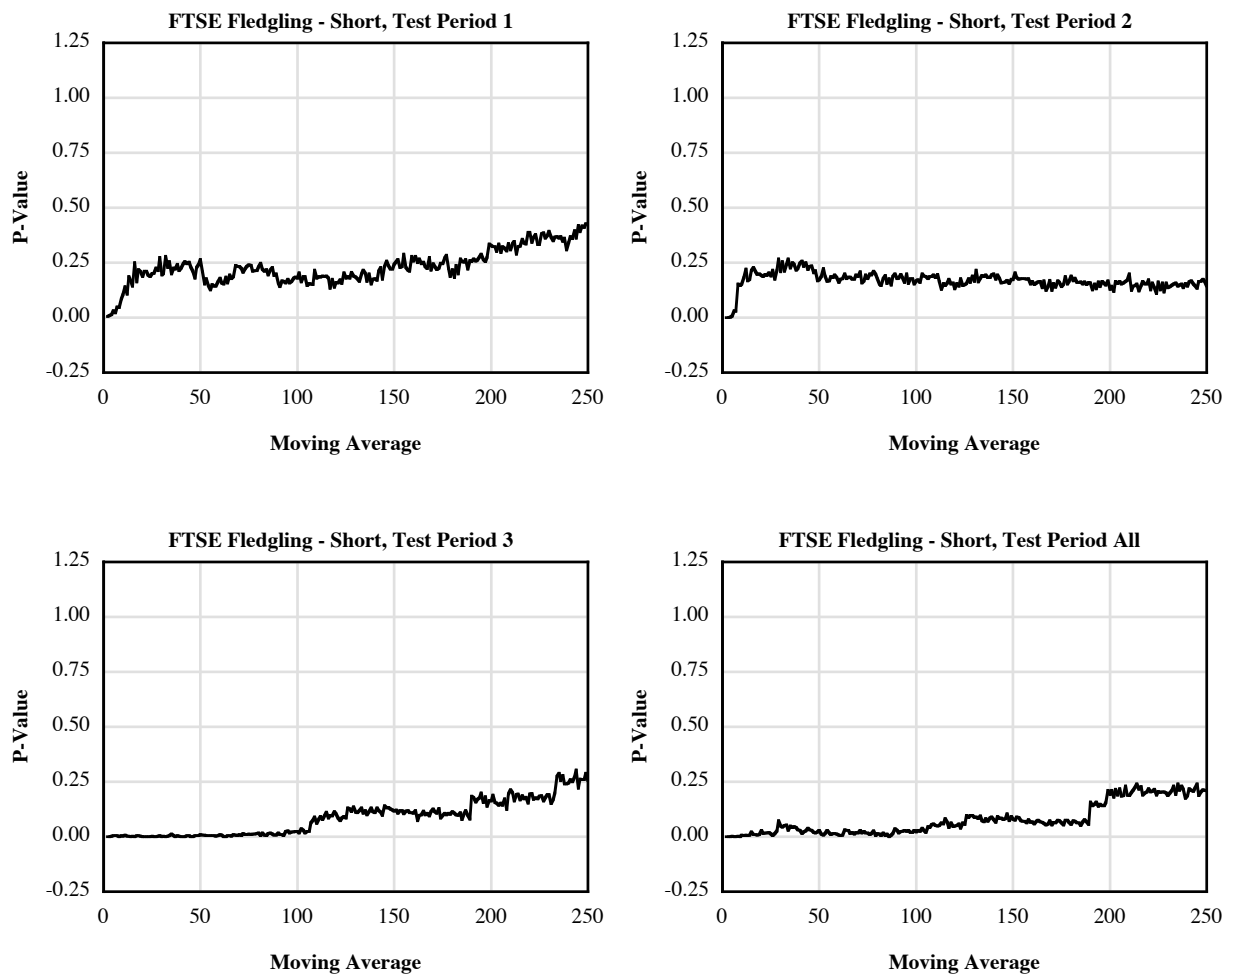


Figure 4.13b Bootstrap p-values for the null hypothesis of no difference in the mean return per trade for the short only price crossover and trade reduction rules for the FTSE Fledgling portfolio.

4.3 Conclusion

A problem with moving average trading, and this applies to technical analysis in general, is whether the buy/sell signals can be said to constitute true information events. At first glance, the answer would appear to be no. Given that the trading rules extract information from past prices and assuming the market is efficient in response to changes in information, the trading rules can only reveal what has happened and not what will happen. True information events precede the buy/sell signals and the trading rules do no more than respond to price changes caused by preceding changes in information or demand.

This is clearly not the case. If the buy/sell signals were not true information events then (1) the return to trading the buy/sell signals should be random in which case the mean return per trade should be zero or negative after costs and (2) the return to buying and selling should be the same. Further, the trade reduction rule shows that not only do the buy/sell signals contain information but also that the sequence in which they occur contains information. Hence, the evolution of the price path contains information. Furthermore, this information is financially exploitable. This is an extremely challenging result. Since the moving average is a smoothed version of the price, the implication is that if this information is visible in the moving average, potentially at least, it is visible for all to see. Assuming trading is the fundamental mechanism by which information is imparted into the price, for trading to be economically viable, it has to be profitable. For trading to be profitable, the information in the price has to be financially exploitable. For the information in the price to be financially exploitable, it has to map to the perception of a trading opportunity where for a trading opportunity to exist, there has to be a feedback mechanism by which the decision to trade can be judged as correct. This feedback mechanism is embedded in the future evolution of the price path. Trading cannot take place without it for the reason that it is otherwise impossible to learn (acquire private information)

about the expected benefits of acting on the information in the price. The trade reduction rule shows that this mechanism exists and is observable in the price path. Once a trade is opened, information about profitability is fed back via the state of the price path and this information is exploited to manipulate the mean return per trade. Chapter 6 expands on this result in more detail.

This has implications for moving average trading as a test of the weak form EMH. An issue with moving average trading is that the trading rules are liable to over trade, which reduces the breakeven costs. The trade reduction rule addresses this issue by reducing the number of losing trades. With reference to Figures 4.6 to 4.9 for example, the mean return per trade for the long only trade reduction rule averaged across all moving averages for all portfolios for the full test period is 84% higher than for the price crossover rule. Assuming this is reflected in a similar increase in the breakeven costs, the trade reduction rule should have more power as a test of the weak form EMH. However, the trade reduction rule is in the market for longer and so takes on more risk. This makes it hard to judge the extent to which the increase in the mean return per trade will be reflected in the breakeven costs after adjusting for risk. A more detailed analysis of this point is further work.

This leaves the question of what information the trading rules are picking up. With reference to Table 3.1 on page 38, excluding the AIM portfolios, the spread increases as company size decreases. The theoretical explanation for this is that market makers are exposed to adverse selection costs. From the market maker's perspective, adverse selection arises when market makers trade with informed traders who have private information about the true value of the stock. Informed traders buy when the price is below its true value and sell when the price is above its true value. Informed trading imparts information into the price whereby informed

buying causes the price to rise and informed selling causes the price to fall. Market makers lose to informed traders since they trade at the wrong price and suffer losses when the price moves against them. The spread allows market makers to recover these losses when trading with uninformed traders. Uninformed traders do not have private information about the true value of the stock and instead trade for reasons other than information. Uninformed trading does not impart information into the price and so uninformed traders pay the spread, which goes to the market makers as profit (Kyle (1985), Glosten and Milgrom (1985), Easley and O'Hara (1987)). Market makers increase the spread to protect themselves against informed trading and thus the spread increases as the probability of informed trading increases.

With reference to Figures 4.6 to 4.13, the mean return per trade also increases as the spread increases. Because informed trading moves the price by virtue of imparting information into the price, this suggests that the trading rules are picking up informed trading. The difference between buying and selling also supports this. One explanation for this difference is that the buy side trades on information whereas the sell side need not. Buy trades indicate an interest in a stock and so are likely to convey firm specific information. Sell trades can be motivated by the need for liquidity, as when selling to raise capital for example, and so need not convey information. Sell trades are also likely to be limited to the range of stocks held in a portfolio and are unlikely to reflect the range of choice available when buying. Buying and selling are asymmetric with the buy side conveying more information than the sell side.

While this argument holds for the FTSE 100 and FTSE 250 portfolios, with the exception of test period 1, it does not hold for the FTSE Small Cap and FTSE Fledgling portfolios where the mean return per trade when selling is greater than the mean return per trade when buying. One explanation is that this is a data artefact attributable to the bias induced by the portfolio

classification scheme and which may or may not disappear when strictly controlling for size. Another explanation is that the information environment is wanting in some way and where this attracts informed traders due to the possibility of an increase in the number of profitable trading opportunities (Grossman and Stiglitz (1980)). If the stocks also trade on low volume, it is likely that informed trading is highly visible either in terms of trade size or according to some other measure of order flow. One reason the mean return per trade is so high is that the portfolios are subject to significant idiosyncratic risk as well as to significant systematic risk in the form of market wide sell-outs due to events such as the technology crash and the sub-prime debt crisis. If informed trading is highly visible during these times and the price falls, there is an incentive to piggyback these trades assuming other traders will see the price fall, note the order flow and also sell. This exposes market makers to unwanted inventory, which reduces prices, which causes more selling, which reduces prices, which causes more selling and so on. The high mean return per trade might be explained by a momentum effect where buying has little impact and selling dominates because everyone is selling. If so, this should be evident in the intra-day order flow data where there should be a lack of buy side interest. Consequently, liquidity should plummet.

Chapter 5

Positive Autocorrelation Rule

This chapter presents the positive autocorrelation rule. Before proceeding, it should be noted that historically, the approach was to first derive the positive autocorrelation rule and to then derive the trade reduction rule at the same time as formalising the test method and analytical framework needed to address the tradability issues raised by the positive autocorrelation rule. The intention being that the positive autocorrelation rule could then be plugged in and tested accordingly. However, the work did not get this far and so it remains in its raw experimental form. First, the data is different. The data is from a different vendor, the test period is shorter and the spread is not available. Second, the test method is different. While the test method of the previous chapter was to fix the capital value of each trade and to test the mean return per trade, the test method in this chapter fixes the trade size and tests the return. The simulation method is also different and assumes a margin account. Questions of statistical significance are not addressed. Last, the positive autocorrelation rule is parametric. Without a method to choose between them, the number of possible trading rules expands considerably. Only one parameter set is tested. The parameters do not follow from a particular method. That said, the positive autocorrelation rule is intended to complement the trade reduction rule and is based on the idea of only trading if it is believed to be profitable to do so. Profitability is treated as equivalent to testing for positive autocorrelation. The explanation for positive autocorrelation in the market microstructure literature is asymmetric information. Informed traders move the

price in the direction of their trades and thereby induce serial correlation. The results in this chapter support this interpretation. For stocks outside the FTSE 100, there is evidence of an information spillover from time t to time $t + 1$ and which appears across a cluster of trading rules, all of which are short term and roughly adjacent in timing. This indicates that positive autocorrelation is prevalent and persistent at this time scale. The positive autocorrelation rule has also been designed to exploit a specific pattern expressed in terms of serial correlation in order flow that occurs in response to evidence of price impact and which persists for several days. However, whether this pattern is solely responsible for inducing the correlation in price changes is not clear. The positive autocorrelation rule can be used as a method to investigate this and other issues at the microstructure level. The remainder of the chapter is organised as follows. Section 5.1 defines the positive autocorrelation rule. Section 5.2 presents the results. Section 5.3 offers conclusions.

5.1 Positive autocorrelation rule

This section defines the positive autocorrelation rule. As stated, the positive autocorrelation rule is based on the idea of only trading if it is believed to be profitable to do so. One way to model this is to think of the trading act as conditional on the presence of a profitable trading opportunity. The intuition behind the positive autocorrelation rule is that the trading act can then be abstracted by modelling the presence of a profitable trading opportunity in terms of positive autocorrelation. Section 5.1.1 motivates the need to test for positive autocorrelation. Section 5.1.2 discusses the variance ratio test and highlights some of its weaknesses. Section 5.1.3 introduces the method of Burgess (2000) as a technique for measuring the deviation of the joint distribution of the variance ratio from the joint distribution for a random walk. The magnitude of the deviation from a random walk is indicative of the extent to which positive autocorrelation is observable in the price. Section 5.1.4 uses this as the basis for defining the

positive autocorrelation rule.

5.1.1 Positive autocorrelation and profitability

A *time series* is a sequence of data points or observations arranged in time order. A common time series in finance is the daily close price $\{p_t\}_{t=1}^T = \{p_1, p_2, p_3, \dots, p_T\}$ where t denotes the time index at which a price is recorded. Define the return r_t as the log return or log price first difference at time t and let r_t have zero mean:

$$r_t = \ln(p_t) - \ln(p_{t-1}) \quad (5.1)$$

$$r_t = r_t - \bar{r}, \quad \bar{r} = \frac{1}{T} \sum_{t=1}^T r_t \quad (5.2)$$

Perhaps the simplest time series model of the price is to assume returns $\{r_t\}_{t=1}^T$ follow a *first order autoregressive* or AR(1) model:

$$r_t = \theta r_{t-1} + e_t \quad (5.3)$$

Where $-1 < \theta < 1$ and e_t is the innovation or error term with mean and variance $(0, \sigma_e^2)$ and which is uncorrelated at all leads and lags. Equation (5.3) states that r_t is a linear function of r_{t-1} and the error e_t . There are three cases to discuss:

1. For $0 < \theta < 1$, r_t exhibits *persistence* in sign. Positive returns tend to be followed by positive returns and negative returns tend to be followed by negative returns. Returns exhibit *positive autocorrelation* in that r_t and r_{t-1} tend to have the same sign and so tend to move in the same direction. This causes the price to trend.
2. For $-1 < \theta < 0$, r_t exhibits *reversion* in sign. Positive returns tend to be followed by negative returns and negative returns tend to be followed by positive returns. Returns exhibit *negative autocorrelation* in that r_t and r_{t-1} tend to have the opposite sign and so tend to move in the opposite direction. This causes the price to mean revert and stops the price from trending.
3. For $\theta = 0$, r_t equals e_t and returns are uncorrelated. This causes the price to follow a *random walk*.

With respect to profitability, moving average trading requires that price changes follow one another in sign. For example, because buying requires the moving average to rise, which in turn requires the price to rise, and because buying is profitable if and only if the sell price is higher than the buy price, which again requires the price to rise, buying requires that rising prices follow rising prices. Similarly, selling requires that falling prices follow falling prices. One way to capture these dynamics is to test for positive autocorrelation. To see this, Figure 5.1 plots three artificial price series simulated according to (5.3) and which exhibit positive autocorrelation, negative autocorrelation and random walk dynamics respectively. The return to trading the price series using a 10-day price crossover rule is also shown. The spread is not simulated and the trading rule is allowed to go both long and short. The return to trading the positive autocorrelation price series is consistently positive, the return to trading the negative autocorrelation price series is consistently negative and the return to trading the random walk

price series is consistently close to zero.¹ In general, assuming the absence of drift, running a moving average trading strategy without testing for positive autocorrelation is likely to result in losses if the required dynamics are not present.

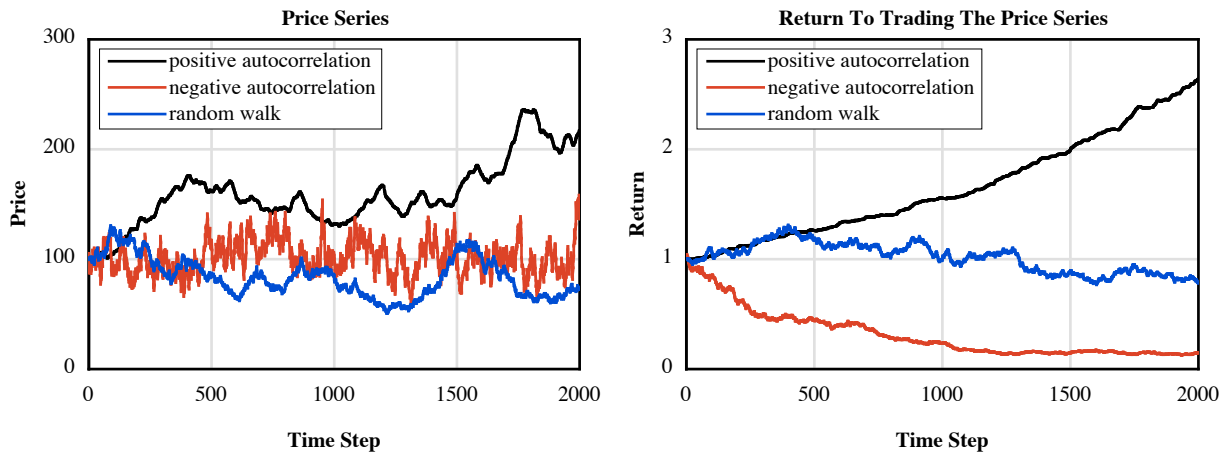


Figure 5.1 Artificial price series that exhibit positive autocorrelation, negative autocorrelation and random walk dynamics and the return to trading the price series using a 10-day price crossover rule.

¹ The return is the final value of each £1's worth of starting capital. A positive return greater than one is profitable and a negative return less than one is loss making.

5.1.2 Variance ratio tests

A common test for serial dependence in time series data is to test whether the autocorrelation coefficients are significantly different from zero. If the coefficients are significantly different from zero, future values of the series are related to previous values and hence the presence of a predictable component in the underlying dynamics. The autocorrelation coefficient $\rho(k)$ at lag k is:

$$\rho(k) = \frac{\sum_{t=k+1}^T (r_t - \bar{r})(r_{t-k} - \bar{r})}{\sum_{t=1}^T (r_t - \bar{r})^2}, \quad \bar{r} = \frac{1}{T} \sum_{t=1}^T r_t \quad (5.4)$$

If $\rho(k)$ is significantly greater than zero, there is evidence of positive autocorrelation at lag k . If $\rho(k)$ is significantly less than zero, there is evidence of negative autocorrelation at lag k . If $\rho(k)$ is not significantly different from zero, there is no evidence of autocorrelation. Popular tests designed to detect departures from zero autocorrelation in either direction at all lags are the Box-Pierce test (Box and Pierce (1970)) and the Ljung-Box test (Ljung and Box (1978)).

A more recent test that is particularly powerful against positive and negative autocorrelation as alternatives to a random walk is the *variance ratio test*. The idea behind the variance ratio test is that the variance of a random walk is linear in the sampling interval. If the variance of the one-period return is σ^2 then (1) the variance of the k -period return should be $k\sigma^2$ and (2) the ratio of the variance of the k -period return to k times the variance of the one-period return should be one. Define the k -period return $r_t(k)$ as:

$$r_t(k) = \ln(p_t) - \ln(p_{t-k}) \quad (5.5a)$$

$$= \ln(p_t) - \ln(p_{t-1}) + \ln(p_{t-1}) - \ln(p_{t-2}) + \cdots + \ln(p_{t-k+1}) - \ln(p_{t-k}) \quad (5.5b)$$

$$= r_t + r_{t-1} + \cdots + r_{t-k+1} \quad (5.5c)$$

The k -period variance ratio $VR(k)$ is the k -period variance divided by k times the one-period variance:

$$VR(k) = \frac{\text{var}(r_t(k))}{k \text{var}(r_t(1))} = \frac{\text{var}(r_t + r_{t-1} + \cdots + r_{t-k+1})}{k \text{var}(r_t)} = 1 \quad (5.6)$$

Lo and MacKinlay (1988) derive asymptotically standard normal test statistics. The variance ratio can be written as:

$$VR(k) = \frac{\frac{1}{Tk} \sum_{t=k}^T (r_t + r_{t-1} + \cdots + r_{t-k+1} - k\bar{r})^2}{\frac{1}{T} \sum_{t=1}^T (r_t - \bar{r})^2}, \quad \bar{r} = \frac{1}{T} \sum_{t=1}^T r_t \quad (5.7)$$

To test for a random walk, Lo and MacKinlay (1988) show that the statistics $Z_1(k)$ and $Z_2(k)$ are distributed as $N(0,1)$ under the null hypotheses of a homoskedastic and a heteroskedastic random walk respectively:²

² For a homoskedastic random walk, the variance of the error term is constant whereas for a heteroskedastic random walk, the variance of the error term changes with time. Estimates of standard errors that assume constant variance are biased by heteroskedasticity and statistical inference is invalid. See, for example, Wright (1980). The $Z_2(k)$ statistic corrects for this.

$$Z_1(k) = \frac{VR(k) - 1}{\sqrt{\phi(k)}}, \quad \phi(k) = \frac{2(2k-1)(k-1)}{3kT} \quad (5.8)$$

$$Z_2(k) = \frac{VR(k) - 1}{\sqrt{\phi^*(k)}}, \quad \phi^*(k) = \sum_{j=1}^{k-1} \left[\frac{2(k-j)}{k} \right]^2 \delta(j), \quad \delta(j) = \frac{\sum_{t=j+1}^T (r_t - \bar{r})^2 (r_{t-j} - \bar{r})^2}{\left[\sum_{t=1}^T (r_t - \bar{r})^2 \right]^2} \quad (5.9)$$

Lo and MacKinlay (1988) also show that the variance ratio is equivalent to a linear function of the first $k-1$ autocorrelation coefficients with linearly declining weights:

$$VR(k) = 1 + \frac{2}{k} \sum_{q=1}^{k-1} (k-q)\rho(q) \quad (5.10a)$$

$$= 1 + \frac{2}{k} [(k-1)\rho(1) + (k-2)\rho(2) + \dots + \rho(k-1)] \quad (5.10b)$$

Hence, variance ratio tests test the null hypothesis $H_0 : \rho(q) = 0 \forall q = 1, 2, \dots, k-1$ against the alternative $H_1 : \rho(q) \neq 0$ for some q . If the variance ratio is significantly greater than one, there is evidence of positive autocorrelation. If the variance ratio is significantly less than one, there is evidence of negative autocorrelation. If the variance ratio is not significantly different from one, there is no evidence of autocorrelation. Figure 5.2 plots the artificial price series of Figure 5.1 along with their variance ratios as defined in (5.7). For the positive autocorrelation price series, the variance increases at a faster rate than the sampling interval and the variance ratio rises above one. For the negative autocorrelation price series, the variance increases at a slower rate than the sampling interval and the variance ratio falls below one. For the random

walk price series, the variance increases linearly with the sampling interval and the variance ratio stays close to one.

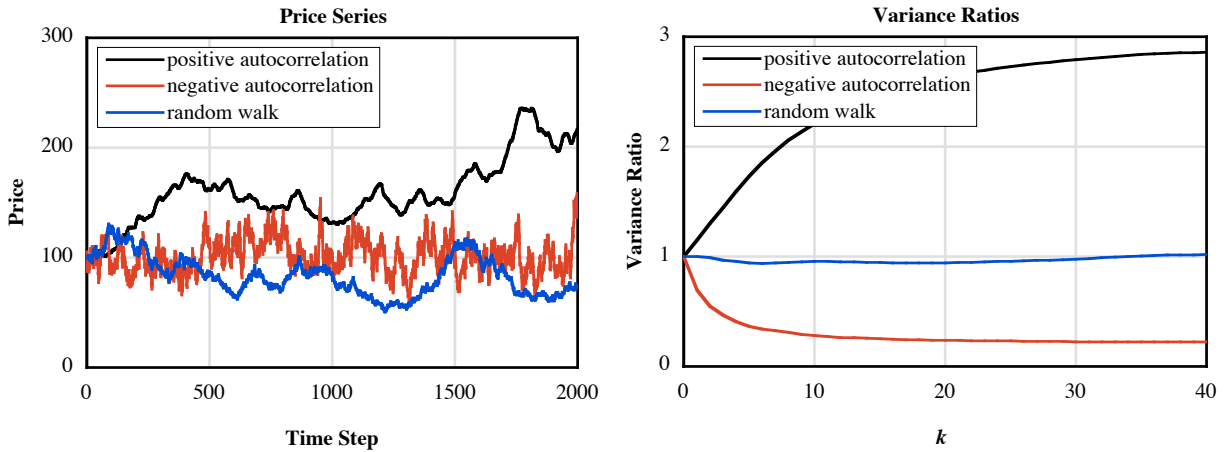


Figure 5.2 Artificial price series that exhibit positive autocorrelation, negative autocorrelation and random walk dynamics and their variance ratios.

There are three problems with variance ratio tests. First, the Lo and MacKinlay (1988) tests are asymptotic tests in that for a finite sample the sampling distributions of the test statistics are approximated by their limiting distributions. Lo and MacKinlay (1989) analyse the finite sample performance of the variance ratio and find that for small samples the null distribution is right skewed and under rejects in the left tail. This can result in misleading inference if the sample size is too small to justify the asymptotic approximations (Cecchetti and Lam (1994)). Wright (2000) extends the Lo and MacKinlay (1988) tests to a non-parametric setting using ranks and signs. This allows the exact simulation of the null distribution for any sample size. Wright (2000) evaluates the rank and sign tests using the same time series benchmarks as Lo and MacKinlay (1989) and concludes firstly that the rank and sign tests are better at detecting

departures from a random walk and, secondly, that they are especially powerful if the returns distribution is non-normal and fat tailed.

Second, it is customary to test several different k values and to reject the null hypothesis if it is rejected for some k . The problem is that the variance ratio test is an individual test whereas testing several different k values is a joint test. Chow and Denning (1993) and Wright (2000) stress that treating individual variance ratio tests as a joint test results in an oversized testing strategy and can lead to over rejection of the null hypothesis. Joint tests that address this test the null hypothesis $H_0 : VR(k_i) = 1 \forall k_i = 1, 2, \dots, k$ against the alternative $H_1 : VR(k_i) \neq 1$ for some k_i . Joint tests include the maximum modulus test of Chow and Denning (1993), the subsampling test of Whang and Kim (2003), the Wald-type test of Chen and Deo (2006) and the bootstrap test of Kim (2006).

Last, given the proliferation of different variance ratio tests, there is little consensus as to the choice of test. It is common in this situation to apply a range of tests with the usual approach being the Lo and MacKinlay (1988) and Wright (2000) tests and one or more joint tests. See, for example, Chang *et al.* (2004), Hoque *et al.* (2007) and Charles and Darne (2009).

5.1.3 Variance ratio profiles

Burgess (2000) treats the joint variance ratio $\mathbf{VP}(k) = [VR(1), VR(2), \dots, VR(k)]$ as having a shape or profile as in the variance ratios of Figure 5.2 and projects this onto the eigenvectors that result from principal component analysis of the covariance matrix of the corresponding profiles for a random walk:³

$$VP(k,i) = (\mathbf{VP}(k) - \overline{\mathbf{RW}}) \mathbf{e}_i \quad (5.11)$$

Where $\overline{\mathbf{RW}}$ is the mean profile for a random walk and \mathbf{e}_i are the eigenvectors or principal components. The magnitudes of the projections onto the eigenvectors are then a measure of the extent to which the profile as a whole deviates from (or is similar to) the archetype for a random walk. This has several advantages. First, (5.11) is a general case test and addresses the problems of the previous section. Second, testing for positive autocorrelation reduces to testing for deviation from the archetype for a random walk. To see this, Figure 5.3 plots the first three eigenvectors for $\mathbf{VP}(100)$ estimated using 1000 random walk time series and the variance ratio defined in (5.7). The magnitude of the projection onto the first eigenvector \mathbf{e}_1

³ Before it is possible to calculate $VP(k,i)$, it is necessary to define $\overline{\mathbf{RW}}$ and \mathbf{e}_i . Both terms are effectively constant. To define them, it is necessary to simulate a number of random walk time series, to calculate $\mathbf{VP}(k)$ for each time series and to save the $\mathbf{VP}(k)$ in the matrix \mathbf{RW} . The mean of \mathbf{RW} is calculated as $\overline{\mathbf{RW}}$. This is subtracted from \mathbf{RW} and principal component analysis is applied to the covariance matrix of \mathbf{RW} to obtain \mathbf{e}_i . The covariance matrix in this case is calculated by the PCA tool. See Jolliffe (2002) for a rigorous introduction to principal component analysis. Note that Burgess (2000) developed the variance ratio profile to test for negative autocorrelation within a statistical arbitrage setting. It is used here to test for positive autocorrelation within a moving average trading setting.

measures the extent to which the profile exhibits positive or negative autocorrelation. For a random walk where each $\mathbf{VP}(k)$ equals one, assuming the mean profile for a random walk $\overline{\mathbf{RW}}$ also equals one, $VP(k,1)$ maps to zero. For positive autocorrelation where each $\mathbf{VP}(k)$ is greater than one, $VP(k,1)$ maps to a number greater than zero. For negative autocorrelation where each $\mathbf{VP}(k)$ is less than one, $VP(k,1)$ maps to a number less than zero. The remaining eigenvectors do not allow this interpretation due to their mixed positive and negative nature. Consequently, to test for positive autocorrelation, it is sufficient to simulate the distribution for $VP(k,1)$ using a random walk and to use this to define a threshold above which deviation from a random walk is deemed significant and thus above which trading can take place. This agrees with Burgess (2000) who performs a Monte Carlo simulation and concludes that as a test for predictability, the first principal component is the best test overall with the remaining principal components better suited to identifying predictable behaviour missed by standard tests including autocorrelation, unit root and variance ratio tests. Third, it is not necessary to limit the archetype to a random walk. For example, it is possible to define (5.11) in terms of previously profitable trades and so (5.11) can be used to test for the similarity to a profitable archetype. This may or may not provide insight into the stock selection problem discussed in Chapter 3. Last, trading rules based on (5.11) are parametric in the variance ratio, the length of data used to calculate the variance ratio, the profile length k , the eigenvector weights, the mean of the random walk archetype and the threshold above which trading takes place. It is possible to imagine these as variables within a genetic algorithm (Goldberg (1989), Mitchell (1998)) and which optimises, say, the risk adjusted mean return per trade measure suggested by Masters (1998). The parametric nature of the trading rules is also a disadvantage however since without optimisation there is no way to choose between them.




Figure 5.3 The first three eigenvectors for $\mathbf{VP}(100)$ estimated using 1000 random walk time series.

5.1.4 Positive autocorrelation rule

Given the discussion of the previous section, for the first principal component, the long only positive autocorrelation rule is:

$$q_1 = 0$$

///// RULE

$$q_{t+1} = q_t$$

IF ($lmin_t$ AND $VP(k,1) \geq x$) THEN $q_{t+1} = 1$

IF ($lmax_t$) THEN $q_{t+1} = 0$ (5.12)

Where $VP(k,1)$ is the projection of the variance ratio profile onto the archetype for a random walk and x is the threshold above which trading takes place. Similarly, the short only positive

autocorrelation rule is:

$$q_1 = 0$$

///// RULE

$$q_{t+1} = q_t$$

IF ($lmin_t$) THEN $q_{t+1} = 0$

IF ($lmax_t$ AND $VP(k,1) \geq x$) THEN $q_{t+1} = -1$ (5.13)

The positive autocorrelation rule introduces two additional problems. Namely, how best to define $VP(k,1)$ and what is a suitable value for x . As discussed previously, because $VP(k,1)$ is highly parametric and because there is no optimisation, the problem of how best to define $VP(k,1)$ is not addressed. Rather, the choice is to use the $R_1(k)$ test of Wright (2000) with a profile length of $k = 25$ calculated over a 130-day moving window. The choice of $R_1(k)$ test is motivated by Wright (2000) who finds that it dominates all other tests in terms of power and does not suffer size serious distortions in the presence of conditional heteroskedasticity. The $R_1(k)$ test is:

$$R_1(k) = \frac{VR(k) - 1}{\sqrt{\phi(k)}}, \quad \phi(k) = \frac{2(2k - 1)(k - 1)}{3kT} \quad (5.14a)$$

$$VR(k) = \frac{\frac{1}{Tk} \sum_{t=k}^T (r_{1t} + r_{1t-1} + \dots + r_{1t-k+1})^2}{\frac{1}{T} \sum_{t=1}^T r_{1t}^2}, \quad r_{1t} = \frac{f(r_t) - \frac{T+1}{2}}{\sqrt{\frac{(T-1)(T+1)}{12}}} \quad (5.14b)$$

Where $f(r_i)$ is the rank of r_i amongst the r_i 's. For $k = 25$, $VP(k,1)$ is then:

$$VP(25,1) = (\mathbf{VP}(25) - \overline{\mathbf{RW}}) \mathbf{e}_1, \quad \mathbf{VP}(25) = [R_1(1), R_1(2), \dots, R_1(25)] \quad (5.15)$$

Figure 5.4 plots the first three eigenvectors for $\mathbf{VP}(25)$ estimated using 1000 random walk time series. The main difference between this and Figure 5.3 is that given a shorter profile length, the weight given to each k is higher. The fact that both sets of eigenvectors exhibit similar shapes indicates that the Burgess (2000) method is robust to the choice of measure. Figure 5.5 plots the variance explained together with the cumulative variance for the first ten principal components. The first principal component explains 85% of the total variance and indicates that not too much information is lost despite limiting the positive autocorrelation rule to the first principal component only.




Figure 5.4 The first three eigenvectors for $\mathbf{VP}(25)$ estimated using 1000 random walk time series.

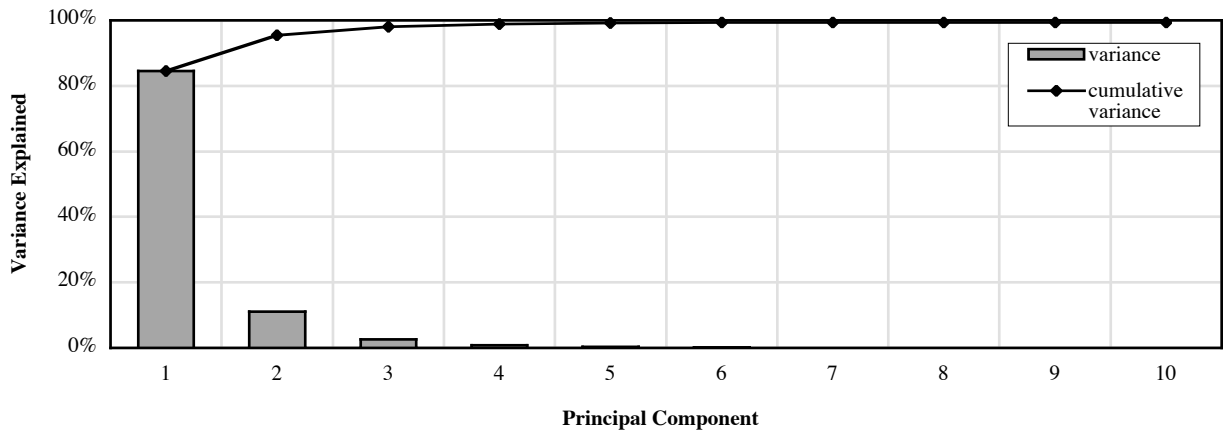


Figure 5.5 Variance explained by the first 10 principal components for **VP(25)**.

Figure 5.6 plots the cumulative frequency distribution for $VP(25,1)$ estimated using another 1000 random walk time series. The 99th percentile occurs at $x = 10$. For $x = 10$ there is a 1% chance of a random walk and for $x > 10$ there is evidence of positive autocorrelation. This is the value used for the threshold x . It should be noted that for $x > 10$ it does not automatically follow that the price is driven by positive autocorrelation alone. It is possible for the price to be driven by a mixture of positive autocorrelation, negative autocorrelation and random walk dynamics but where the positive autocorrelation term is sufficiently dominant for $VP(25,1)$ to be greater than 10. All that can be said in this situation is that while a positive autocorrelation term is likely to exist and so is likely to have an impact, the exact form of the price dynamics remains unknown.




Figure 5.6 Cumulative frequency distribution for $VP(25,1)$ estimated using another 1000 random walk time series.

Despite the intuition of combining moving average trading with a variance ratio test, it is not known of any work in this area. The only work known to apply the method of Burgess (2000) is Lindemann *et al.* (2005) who leave out the subtraction of the mean in the variance ratio of (5.7) in order to allow for a random walk with drift. The first two principal components are then plotted in (x, y) space to give a 2-dimensional structure visualisation tool. Trading rules are not discussed. The only work known to explicitly combine moving average trading with an autocorrelation model in a way similar to the positive autocorrelation rule is Fang and Xu (2003) who estimate a rolling AR(1) model for the Dow Jones Industrial, Transportation and Utilities indices. The model forecasts are combined with moving average trading rules such that trading only takes place when both the AR(1) model and the trading rule emit a buy/sell signal. This doubles the breakeven costs relative to the moving average trading rules on their own. However, time variation in the model coefficient is not discussed and so the exact form of the AR(1) models is not known. Nevertheless, this is an interesting result since Brock *et al.*

(1992) find that the returns to moving average trading are not explained by an AR(1) model although they do not consider locally fitted models.

The positive autocorrelation rule can also be thought of as a nonlinear model that switches trading on/off whenever the price enters/exits a positive autocorrelation regime. Brock *et al.* (1992) suggest that the difference in the returns uncovered by moving average trading might be explained by nonlinear models with asymmetric dynamics. Gencay (1998) finds evidence of nonlinear predictability when the buy/sell signals from moving average trading rules are used as inputs to a single-layer feedforward neural network. Gencay and Stengos (1998) find that predictability increases significantly when volume is used as an additional input. Self and Mathur (2006) estimate a momentum threshold autoregressive (MTAR) model and use this together with the asymmetric stationarity test of Enders and Granger (1998) to identify non-stationary, symmetric stationary and asymmetric stationary regimes in the G7 national stock indices. There is evidence that the return to moving average trading varies with each regime, particularly so with regard to volatility, although there is no attempt to design an appropriate trading strategy. Choe *et al.* (2011) estimate nonlinear autoregressive (NAR) models for the G7 national stock indices. Following Nam *et al.* (2005), trading rules that exploit patterns in previous returns are found to be significant but only when nonlinearity results in a consistent asymmetric pattern in the model coefficients. Chong and Lam (2010) estimate rolling AR(1) and self-exciting threshold autoregressive (SETAR) models for the DJIA, NASDAQ, NYSE and S&P 500 indices. The model forecasts generate higher returns than comparable moving average trading rules in all cases. However, this is before costs. It is not clear if the models generate more or less trades than the moving average rules.

It should also be noted that there is a literature that tests the weak form EMH based on tests applied within rolling windows. Examples include rolling variance ratio tests (Tabak (2003), Kim and Shamsuddin (2008)), rolling unit root tests (Phengpis (2006)), rolling bicorrelation tests (Lim (2007)), rolling Hurst exponents (Cajueiro and Tabak (2006)) and rolling AR(1) models (Lo (2004), Ito and Sugiyama (2009)). The general result is that there is evidence of time varying deviation from a random walk. This suggests that market efficiency is not an all or nothing condition and that there are times when there is evidence of inefficiency.

5.2 Results

This section presents the results for the positive autocorrelation rule. Section 5.2.1 describes the data. Section 5.2.2 discusses the simulation method. The simulation method assumes an execution only margin account and fixes the trade size. Section 5.2.3 presents an illustrative example where it is shown that the positive autocorrelation rule is profitable and that profits appears across a cluster of trading rules, all of which are short term and roughly adjacent in timing. This captures the intuition of positive autocorrelation as a short term dynamic and is evidence that positive autocorrelation is prevalent and persistent at this time scale. Section 5.2.4 presents the results for the FTSE 100, FTSE 250, FTSE Small Cap and FTSE Fledgling portfolios. The results are the same. This prompts a discussion of the strategy underlying the positive autocorrelation rule and it is suggested that this should be observable in the adverse selection component of spread decomposition models. The positive autocorrelation rule can be used to investigate this at the microstructure level. There is also evidence of price impact.

5.2.1 Data

As stated in the introduction, due to data limitations, the test data is different to that used for the trade reduction rule. The data comprises 13 years of daily close prices for stocks listed on the London Stock Exchange from 06-January-1994 to 06-January-2007. The data is from the Sharescope and Updata Technical Analyst software packages, both of which are supplied with historical price databases. The data comes as a sequence of date-stamped records of the form $(date_t, close_t)$ where $date_t$ is the date for which the price is available and $close_t$ is the end of day close price. Prices are adjusted for stock splits and other capital events but do not include the dividend. Prices exclude weekends but include public holidays. Prices for public holidays are duplicates of the previous days prices. Prices for public holidays are not treated differently and are retained along with other prices.

The first step in pre-processing is to check that the price is available for the whole of the test period. If not, the price file is deleted. The next step is to check that the prices from the two data sources equal one another. If not, the price file is deleted. This gives 253 stocks in total. Investment trusts are not included. Once the data is pre-processed, the stocks are sorted into portfolios. There is one portfolio for each of the FTSE 100, FTSE 250, FTSE Small Cap and FTSE Fledgling indices. If a stock is a member of the FTSE 100 index, it is allocated to the FTSE 100 portfolio. If it is a member of the FTSE 250 index, it is allocated to the FTSE 250 portfolio and so on. The indices are supplied by the FTSE Group (www.ftse.com). The date at which the indices are correct is not available. A complete listing of each portfolio is given in the appendix. There are 43 stocks in the FTSE 100 portfolio, 77 in the FTSE 250 portfolio, 74 in the FTSE Small Cap portfolio and 59 in the FTSE Fledgling portfolio.

5.2.2 Simulation method

Simulation assumes a retail client, execution only margin account. Such accounts are readily available to the public and where the underlying financial instrument is either a contract for difference or a spread bet. Margin allows for taking a position without needing the capital to fund the full value of the position. For example, going long or short £10,000 worth of shares with margin of 10% requires initial capital of $£10,000 \times 10\% = £1,000$. The effect is to gear the return by reducing the capital commitment. Margin needs to be financed however. Long positions pay interest on the overnight value of the position whereas short positions receive interest on the overnight value of the position. The interest rates used to simulate the trading rules are 7.5% for long positions and 2.5% for short positions. For example, the interest paid on a long position with an overnight value of £10,000 is $£10,000 \times (7.5\% / 365.25) = £2.05$ and which is debited from the account at the start of the next day. The interest received on a short position with an overnight value of £10,000 is $£10,000 \times (2.5\% / 365.25) = £0.68$ and which is credited to the account at the start of the next day.

For simulation purposes, the margin requirement is 100%. There has to be sufficient capital available to fund the full value of each trade. The reason for this is to determine whether the positive autocorrelation rule has value as a real world decision support tool assuming trading is leveraged and is in units of one normal market size. To simulate the costs associated with leveraging the trading rules, they are subject to financing. It is then possible to determine if the trading rules are stable in the presence of financing independently of gearing the return. To simulate trading in units of one normal market size, the trade size is fixed at 1000 shares. The problem with this is that because the trading rules are not fully invested, it is not known how much capital is needed to fund trading until the end of the simulation. Similarly, fixing

the trade size introduces an element of self-financing in that the profits from previous trades can be used to pay for future trades. The approach in this case is to first simulate the trading rules and to calculate the cumulative cash profit. Once the cumulative cash profit is known, the minimum capital needed to guarantee that each trade is executed is added in. This means that the return is calculated *ex-post* and is the maximum return achievable for the minimum capital needed to fund trading to the end of the simulation. This is calculated on a trading rule by trading rule basis and is different for each trading rule for each moving average. While this might appear to bias the return, this is not the intention. Rather, it is to highlight evidence of profitability. For example, previewing the results, compared to the price crossover rule, the positive autocorrelation rule needs less capital to fund trading. Given that uninvested capital does not earn interest, equalising the capital across the trading rules suppresses the return to the positive autocorrelation rule due to the high proportion of uninvested capital. Calculating the return as the maximum achievable for the minimum capital needed to fund trading does not suppress the return and instead highlights evidence of profitability. This is true for both the price crossover rule and the positive autocorrelation rule.

Because the spread is not available, round trip transaction costs are levied at 2.5% per trade. For portfolios, the minimum capital needed to fund trading is calculated for the whole of the portfolio. Stocks within the portfolio are traded individually and the minimum capital needed to fund trading is the minimum capital needed to trade all of the stocks at the same time. This includes margin calls. The return to a trading rule is then a cash equity curve that starts from the minimum capital. Uninvested capital does not earn interest. Dividends are not simulated. The moving average is as described in Section 3.2. Table 5.1 lists the parameters used for the simulations.

Table 5.1

Simulation parameters

Margin	Long Interest	Short Interest	Trade Size	Round Trip Costs
100%	7.5%	2.5%	1000 shares	2.5%

Margin of 100% means that the return is subject to financing but is not geared.

5.2.3 Illustrative example

Figure 5.7 plots the equity curves using a 10-day moving average for the price crossover and positive autocorrelation rules applied to Yule Catto (a FTSE Small Cap chemicals company) for the period 06-January-1994 to 06-January-2007. The simulation parameters are listed in Table 5.1. Note that long and short trades are not treated separately. The trading rules can go both long and short. The price crossover rule is:

$$q_1 = 0$$

///// RULE

$$q_{t+1} = q_t$$

$$\text{IF } (lmin_t) \text{ THEN } q_{t+1} = 1$$

$$\text{IF } (lmax_t) \text{ THEN } q_{t+1} = -1 \tag{5.16}$$

The positive autocorrelation rule is:

$$\begin{aligned} q_1 &= 0 \\ \text{///// RULE} \\ q_{t+1} &= q_t \\ \text{IF } (lmin_t) \text{ THEN } q_{t+1} &= 0 \\ \text{IF } (lmax_t) \text{ THEN } q_{t+1} &= 0 \\ \text{IF } (lmin_t \text{ AND } VP(25,1) \geq 10) \text{ THEN } q_{t+1} &= 1 \\ \text{IF } (lmax_t \text{ AND } VP(25,1) \geq 10) \text{ THEN } q_{t+1} &= -1 \end{aligned} \tag{5.17}$$

For the price crossover rule, the equity curve rises from January-1994 until September-2002 indicating that the price crossover rule is profitable and that the price is trending. The equity curve then falls after September-2002 indicating that the price crossover rule is loss making and that the price is no longer trending. The change in the price dynamics causes all previous gains to disappear resulting in a net loss. For the positive autocorrelation rule, trading simply switches off after September-2002 resulting in a net profit. In terms of performance, the price crossover rule generates 298 trades and has a win rate of 24% with a mean return per trade of -0.4%. The positive autocorrelation rule generates 77 trades and has a win rate of 43% with a mean return per trade of 2.6%. The price crossover rule also needs 3 times more capital than the positive autocorrelation rule.




Figure 5.7 Equity curves for the 10-day moving average price crossover and positive autocorrelation rules for Yule Catto.

Figure 5.8 plots the return for moving averages in the range $n = 2, 3, \dots, 100$. The return to the positive autocorrelation rule is consistently higher than for the price crossover rule. The positive autocorrelation rule is also profitable for a range of moving averages,⁴ all of which cluster together and capture the intuition of positive autocorrelation as a short term dynamic. Although an isolated example, this is a fundamental result. As the moving average increases, so the lag of the trading rules increases. As the lag of the trading rules increases, so the time delay in the buy/sell signals increases. Consequently, there is little to separate neighbouring trading rules other than the delay in the timing of the buy/sell signals. Because profits appear

⁴ From the data in Chapter 3, the mean spread for Yule Catto for the period 06-January-1994 to 06-January-2007 is 2.1%. Round trip costs transaction costs of 2.5% are therefore likely to overestimate the spread. The positive autocorrelation rule is also subject to financing. This is an additional cost. Conclusions regarding profitability are robust in this respect.

across a cluster of trading rules all of which are roughly adjacent in timing, there is evidence that positive autocorrelation is both prevalent and persistent at this time scale. This takes the form of an information spillover from time t to time $t + 1$ and which is sufficiently prevalent for the positive autocorrelation rule to profit from it and which is sufficiently persistent for it appear across a cluster of trading rules.




Figure 5.8 Return for the price crossover and positive autocorrelation rules for Yule Catto.

5.2.4 Portfolio simulations

Figure 5.9 plots the results for the price crossover and positive autocorrelation rules for the FTSE 100, FTSE 250, FTSE Small Cap and FTSE Fledgling portfolios. The trading rules are the same as in the previous section and are defined by (5.16) and (5.17). Moving averages are in the range $n = 2, 3, \dots, 100$. The return for the positive autocorrelation rule is consistently higher than for the price crossover rule.⁵ However, with regard to profitability, the results are difficult to interpret. This is because it is not clear to what extent the results are robust to the spread. For example, with reference to Table 3.1 on page 38, round trip transaction costs of 2.5% are likely to overestimate the spread for the FTSE 100 and FTSE 250 portfolios and to underestimate the spread for the FTSE Small Cap and FTSE Fledgling portfolios. This limits discussion of profitability to the FTSE 100 and FTSE 250 portfolios. Nevertheless, in general, the positive autocorrelation rule is able to identify trades that are more profitable than for the price crossover rule. For the FTSE 250 portfolio, these trades are also financially exploitable. Although limited, this result implies that profitable trading opportunities are observable in the price.

⁵ The apparent discrepancy in the results for the price crossover rule for this chapter and the previous chapter is explained by the difference in strategy. The strategy for this chapter is to fix the trade size. This does not break up serial dependence in the order of the trades and so losses are compounded. Trades are also subject to financing which further reduces the return. The strategy in the previous chapter was to fix the capital value of each trade. This breaks up serial dependence in the order of the trades so that each trade earns the mean return per trade by definition. Trades are also not financed and so the return is not reduced. Because it breaks up serial dependence in the order of the trades, the better strategy is to fix the capital value of each trade.

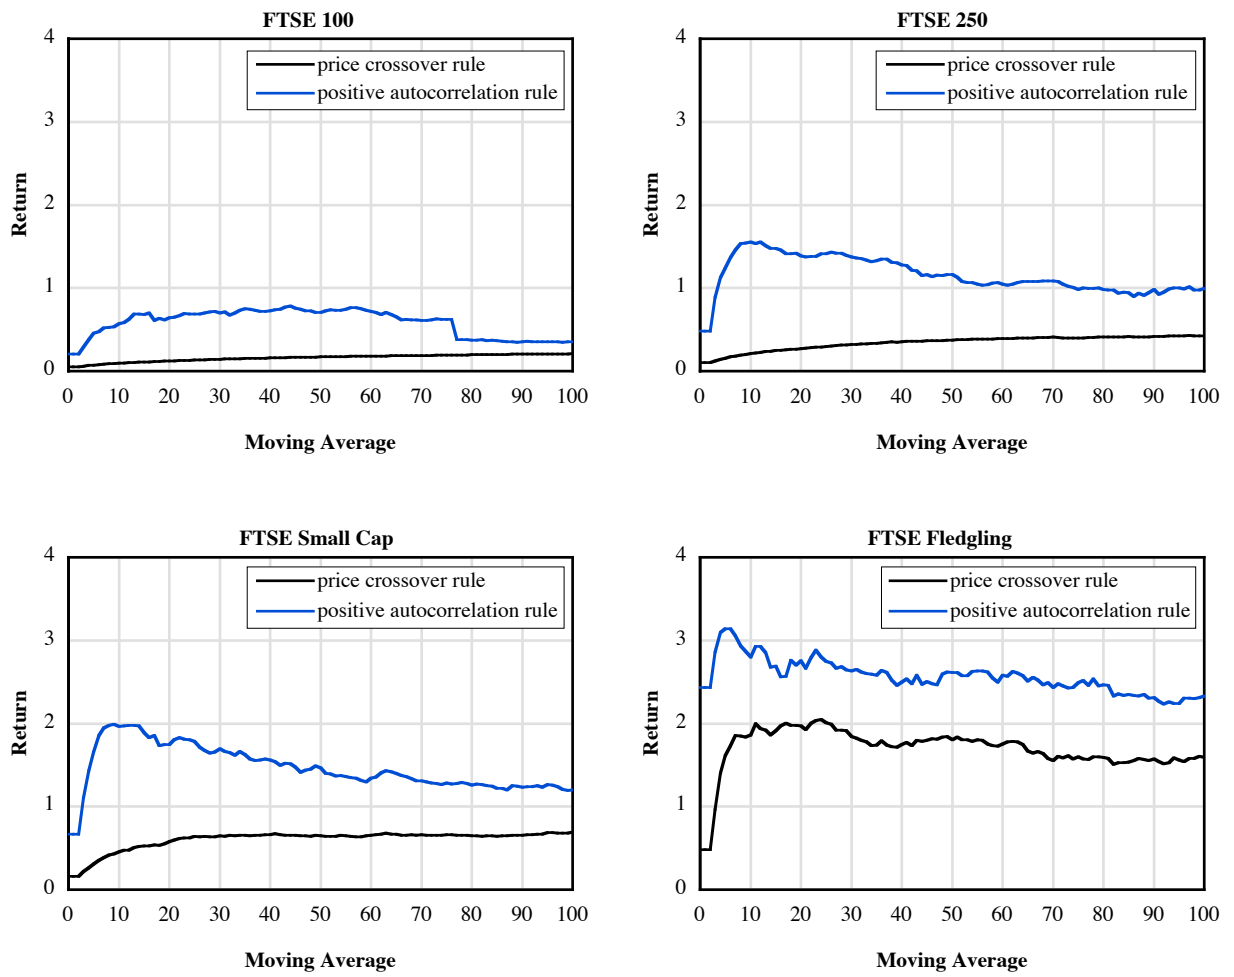


Figure 5.9 Return for the price crossover and positive autocorrelation rules for the FTSE 100, FTSE 250, FTSE Small Cap and FTSE Fledgling portfolios.

It is also noticeable that the return increases as company size decreases. This suggests that as with the trade reduction rule, the positive autocorrelation rule is picking up informed trading. This is a significant result. The positive autocorrelation rule is designed to exploit a specific pattern. For example, suppose the price is p_t and suppose the trader believes the correct price is p_t^* where $|p_t^* - p_t| > c$. This means that p_t is indicative of a profitable trading opportunity. The optimal strategy is to trade p_t and for the price to immediately change to p_t^* whereupon a profit is realised. However, for the trader, it is unrealistic to expect the price to immediately change to p_t^* in response to a single trade. Suppose then that the trader knows other traders will also recognise p_t as a profitable trading opportunity and so instead of trading, waits for these other traders to appear in the order flow. Once these traders start to appear in the order flow, if there is evidence that p_t is sensitive to their trades in that p_t starts to move toward p_t^* , the trader also trades. Suppose there are many traders following this strategy. They also respond to the sensitivity of the price to the order flow and trade as well. This induces serial correlation in the order flow. Trades tend to be in the same direction and the price changes accordingly. If this persists for several days, there is an information spillover from time t to time $t + 1$ and the positive autocorrelation rule is profitable. The positive autocorrelation rule is specifically designed to exploit this pattern by searching out times when it is likely to hold. With reference to Figure 5.9, not only is there evidence to suggest that it exists but also that it becomes increasingly visible as company size decreases.

The significance of this is that the trading patterns and/or information dynamics underlying the positive autocorrelation rule should be observable in the intra-day order flow data. More specifically, assuming the positive autocorrelation rule is profitable, it should identify times when there is variation in adverse selection costs. Popular spread decomposition models that

measure the adverse selection component of the spread include Glosten and Harris (1988), George *et al.* (1991), Lin *et al.* (1995), Huang and Stoll (1997) and Madhavan *et al.* (1997). However, a general issue with these models is that they often result in different estimates of the adverse selection component and which are sometimes implausible (Neal and Wheatley (1998), Van Ness *et al.* (2001), Chelley-Steeley and Park (2008)). A potential application of the positive autocorrelation rule is to use it to investigate the adverse selection component of the spread by splitting the price into times when the positive autocorrelation rule is profitable, not profitable and not trading. Estimates of adverse selection should vary accordingly. If not, either the adverse selection component of the models is misspecified or there is an unknown source of asymmetric information. It is not known of any work that explicitly links moving average trading and adverse selection in this way.

Similarly, apart from the FTSE 100 portfolio, the return for the positive autocorrelation rule exhibits a distinct hump shape. Returns are maximised in the short term and decrease as the length of the moving average increases. This implies that the information uncovered by the positive autocorrelation rule is short lived. Empirical studies such as Hasbrouck (1991) and Easley *et al.* (1997) suggest that informed traders with short lived information will prefer to trade large trade sizes. If so, variation in trade size should also be observable in the intra-day order flow data and the positive autocorrelation rule can be used to investigate this. It is also worth noting that in London, market makers are only obliged to quote prices good for orders up to one normal market size. If informed traders do prefer to trade large sizes, a measure of informed trading might be the price impact of trades normalised by the normal market size.⁶

⁶ Price impact refers to the correlation between a trade and the subsequent price change. See Hasbrouck (2007).

This measures the sensitivity of quote revisions to trades in units of normal market size. The sensitivity of quote revisions to trades in units of normal market size is thought to be a trigger underlying the positive autocorrelation rule.

There is also evidence of price impact. The fact that the price crossover rule is profitable for the FTSE Fledgling portfolio indicates that prices exhibit a degree of persistence. Additional evidence can be found by looking at the ratio of the capital needed to fund the trading rules. For moving average n , let $capital(pc)_n$ denote the capital needed to fund the price crossover rule and let $capital(pa)_n$ denote the capital needed to fund the positive autocorrelation rule. Figure 5.10 plots the ratio of the capital needed to fund the price crossover rule relative to the capital needed to fund the positive autocorrelation rule:

$$ratio = \frac{\sum_{i=2}^n capital(pc)_i}{\sum_{i=2}^n capital(pa)_i} \quad (5.18)$$

As company size decreases, the capital needed to fund the price crossover rule tends to that needed to fund the positive autocorrelation rule. This cannot be explained by volatility since volatility increases as company size decreases. This implies that as company size decreases, prices exhibit greater persistence. Trades are more likely to have a permanent impact, which is likely to be a contributory factor in the return for the positive autocorrelation rule.




Figure 5.10 Ratio of the capital needed to fund trading for the price crossover and positive autocorrelation rules.

5.3 Conclusion

The main result to come from this chapter is that in general, the positive autocorrelation rule shows that profitable trading opportunities are observable in the price. For stocks outside the FTSE 100, there is evidence of an information spillover from time t to time $t + 1$. The hump shape seen in the return indicates that this information is short lived. It also indicates that this information is sufficiently prevalent for the positive autocorrelation rule to profit from it and sufficiently persistent for it to appear across a cluster of trading rules, all of which are roughly adjacent in timing. However, evidence of profitability is limited to the FTSE 250 portfolio. It is not clear to what extent the results for the FTSE Small Cap and FTSE Fledgling portfolios are robust to the spread. It is also not clear whether simulating long and short trades together adds to or takes away from the return. Similarly, with reference to Figure 5.8, it is not clear if trading is continuous for the whole of the test period. To draw stronger conclusions regarding

profitability it is necessary to simulate the positive autocorrelation rule using the data and test method of the trade reduction rule as well as to investigate whether trading is continuous.

It is also noticeable that the return increases as company size decreases. This suggests that as with the trade reduction rule, the positive autocorrelation rule is picking up informed trading. However, underlying informed trading is the argument that it enhances market efficiency by making prices more informative. This is in conflict with the evidence that the price dynamics uncovered by the positive autocorrelation rule are profitable. The trading patterns underlying the positive autocorrelation rule are therefore of great interest. A contribution of the positive autocorrelation rule is to provide a method by which to peel back the price and to investigate these issues at the microstructure level. Of particular note is that the positive autocorrelation rule has been designed to exploit a specific pattern. This is expressed as serial correlation in order flow that occurs in response to evidence of price impact and which persists for several days. However, whether this pattern is solely responsible for inducing the correlation in price changes is not known. Biais *et al.* (1995) suggest that serial correlation in order flow can be explained by (1) the splitting of large orders, (2) piggybacking whereby traders follow what other traders are doing and (3) similar reactions to the same events. Which of these is able to explain the correlation in price changes should be observable in the intra-day order flow data. The same argument holds for spread decomposition models. The positive autocorrelation rule splits the price into times when it is profitable, not profitable and not trading. Estimates of the model parameters should vary accordingly. To summarise, the positive autocorrelation rule is clearly able to identify times when the price behaves in a way that is different from the norm. The reasons for this difference should be observable in the intra-day order flow data. A more detailed analysis of the positive autocorrelation rule at this level is further work.

As discussed in Chapter 3, moving average trading suffers from a stock selection problem. If the price dynamics exploited by moving average trading are not in the price, the trading rules will not find them. A potential solution to this problem is to record $VP(k,i)$ at each time step whereby each stock has a $VP(k,i)$ signature. Ranking each stock at each time step according to some function of its signature is a method by which to solve the stock selection problem. Given their ranking, it is then possible to conduct a market wide search for those stocks most suited to moving average trading. The maximally ranked stocks might be those with maximal evidence of positive autocorrelation for example. There are two points. First, if the stocks are ranked and sorted into portfolios, if the results for the positive autocorrelation rule vary with the portfolios, there is the basis for a deeper investigation into the reasons for this. Second, it is important to be aware that the positive autocorrelation rule is a functional decision support tool. Ranking each stock at each time step allows the causes for their ranking to be assessed. If appropriate, this information can be traded.

Chapter 6

Conclusion and Further Work

The question asked in this thesis is whether remodelling moving average trading to reduce the number of losing trades increases the mean return per trade to the extent that the trading rules are profitable and, if so, whether this is economically significant. This is motivated by market efficiency. A general issue with market efficiency is that it is extremely difficult to rationalise trading if markets are efficient as defined by Fama (1970). If the price impounds all available information, other than in response to new information, why is it that traders trade? There are four possibilities. First, markets are efficient in the sense of Jensen (1978) and Taylor (2005). Profitable trading opportunities exist but where it is not possible for traders to outperform the market on a risk adjusted basis. However, traders might leverage their trades and outperform the market if measured by the return on capital employed. Second, markets are inefficient in the sense of Grossman and Stiglitz (1980). It is possible for traders to outperform the market but where outperformance is limited by the cost of acquiring information. Third, markets are inefficient. Profitable trading opportunities abound and outperformance is unlimited. Fourth, all of the above hold in some way. Market efficiency as defined is not rich enough a concept to capture the complexities of the market and/or the trading process.

To examine these issues further, the approach is to adopt the perspective of the trader. Two trading rules are proposed, both of which are designed to capture trading practice. The trade

reduction rule is based on the idea of allowing a trade to run and the positive autocorrelation rule is based on the idea of only trading if it is believed to be profitable to do so. The trading rules follow from an understanding of what it is to trade, are independent of the data and are supported mathematically. The problem with modelling trading practice is that it implies the need to model the trading decision. To capture this, the method is to define trading as driven by price direction and to replace the information that would normally be input to the trading decision with the information contained in the moving average buy/sell signals. This allows the trading decision to be modelled in terms of buy/sell actions that transform the underlying trade distribution. The advantage is that the modelling problem reduces to describing how a trader might exploit price direction and where the moving average renders this testable. The disadvantage is that the moving average is a crude tool. However, this crudity is intentional. What is being modelled is not the price but the trader's *response* to the price. If it is possible for the trader to exploit the information in the moving average, not only is there evidence of market inefficiency but given that the moving average is little more than a smoothed version of the price, there is also evidence that this is observable in the price. The implication being that if market inefficiency is observable in the price, potentially at least, it is visible for all to see.

The easiest way to put this in context is to consider the following scenario. Suppose a trader opens a trade. Suppose also that this is solely for the purposes of making a profit. For this to occur, the price needs to convey two pieces of information. First, it needs to convey that the price represents a trading opportunity. Second, it needs to convey that in the event of acting on a trading opportunity, all else being equal, a profit will be realised. Figure 6.1 illustrates this in terms of the trader's perception of the relationship in information flow.



Figure 6.1 Trading and the trader's perception of the relationship in information flow.

Note that Figure 6.1 is independent of the actual trading strategy. It is sufficiently general to capture all trading strategies irrespective of their origin, complexity or implementation. The first contribution of this thesis is that the trading rules recover the model of Figure 6.1. The trade reduction rule relates to the feedback on the decision to trade. After a trade is opened, information on profitability is fed back via the state of the price path and this information is used to manipulate the mean return per trade. The positive autocorrelation rule relates to the profitability of a trading opportunity. Before a trade is opened, profitability is evaluated and this information is used to manipulate the return. Both trading rules uncover information that is missed by the price crossover rule. The information uncovered by the trade reduction rule is long lived and the information uncovered by the positive autocorrelation rule is short lived.

The second contribution of this thesis is to show that the information uncovered by the trade reduction and positive autocorrelation rules is financially exploitable. At its most basic level, this means that the information needed for trading to be economically viable is observable in

the price. A general issue with trading is whether it is elitist. Can the average trader expect to profit or is profitability the preserve of the select few? The choice of moving average trading addresses this issue and there is evidence that the information underlying Figure 6.1 is visible for all to see. This implies that first, for the test data at least, the market cannot be efficient in the sense of Fama (1970). The results do not support this perspective. Rather, second, there is support for learning. At the microstructure level, this implies an explanation for uninformed traders. A problem with uninformed traders is why they continue to trade if they continue to lose (O'Hara (2003)). Uninformed traders can be defined as attempting to learn the model of Figure 6.1. If successful, they become informed traders. If unsuccessful, they can be expected to stop trading and to leave the market. This supports the adaptive market hypothesis (AMH) of Lo (2004) where markets comprise continuously coevolving competing trading strategies. Third, given the support for learning, it should be possible to acquire expertise. This implies that there should exist traders who consistently generate profits. There is some evidence that such traders exist. Puckett and Yan (2011) find that the trades of institutional fund managers generate an average excess of 20 to 26 basis points per annum after costs. While this appears to be a small number, for average net assets of US\$22 billion, this is a large number in cash terms. When ranked by previous trading performance, for the top quintile, there is evidence that performance persists from quarter to quarter. Barber *et al.* (2011) find that the trades of the top ranked day traders in Taiwan generate an average excess of 28.1 basis points per day after costs. At the portfolio level, this reduces to 2 basis points per day after costs. As before, previous trading performance is the best predictor of future performance. Neither study finds evidence to support the hypotheses that returns are explained by liquidity provision or inside information. Results such as these are important and indicate that to fully understand market efficiency it is necessary to study the individual trading records of skilled traders. This thesis is a step in that direction. Fourth, there is support for technical analysis. Technical analysis is

used to visualise information. What is required of the trader is to decide if this information is meaningful. Last, there is support for trading system design. The trade reduction and positive autocorrelation rules are functional trading tools. However, the theme to the results is that by modelling elements of decision making as found in the real world it is possible to show that the market information environment is financially exploitable. This introduces an important architectural issue. The majority of trading systems automate the trading process and design the trader out. A different approach and which allows greater variation in decision-making is to concentrate on decision support tools and to design the trader in.

The third contribution of this thesis is the relationship with market microstructure. A slightly different way to think of Figure 6.1 is that market microstructure looks out of the market and in the direction of the trader. The approach in this thesis is to look out of the trader and in the direction of the market. In doing so, the trading rules are the glue that binds the two together. The strongest interpretation that can be applied to the trading rules in this context is that they are in fact examples of informed trading. Compared to the price crossover rule, they are more able to extract meaning from (or more able to understand) the same price information. Given this interpretation, it is perhaps not surprising to find that the trading rules uncover issues of informed trading, liquidity, adverse selection and price impact. Intuitively, this is correct. Of particular note is that the positive autocorrelation rule is a method by which to peel back the price and to investigate these issues using high frequency data. The positive autocorrelation rule has the potential to be a fruitful test bed in this respect. It is not known of any work that explicitly links moving average trading and market microstructure in this way.

This leaves the question of economic significance. This is a difficult question and one that it is not possible to address without further testing. However, there is evidence that the market

information environment is financially exploitable. It is reasonable to assume that traders are aware of this and are able to profit from it. A testable prediction of this thesis is that the high frequency data should contain evidence of traders who consistently generate profits. Whether this also implies economic significance is not clear however. The issue with data of this type is that each identifier needs to relate to a single trader and where the data needs to include all of the trader's activity. Clearly, there will always be some doubt about this. It is also the case that the trader's capital is an unobservable variable. This is perhaps the bigger problem since the trader's capital is the variable of interest. The advantage of trading rules is that economic significance is always testable. The disadvantage is that by reducing trading to an algorithm, they generally fail to capture the broader complexity of the trading process. The work in this thesis is intended to address this and the results suggest that economic significance is not by definition unobtainable. It is within reason to hypothesise that this is a fundamental property of financial markets and is one that is necessary for them to exist given that it is the promise of economic significance, or at least the perception of this promise that attracts traders.

6.1 Further work

It is thought that the market efficiency test most likely to generalise in the sense of also being suitable for large cap stocks is to extend the trade reduction rule to include an autocorrelation test but where trading switches off in the presence of negative autocorrelation. This needs to be investigated. The trading rules can also be extended to include the location of the buy/sell signals within $[-1, 1]$. For example, for $-1 \leq y, z \leq 1$, this gives trading rules of the type:

$$\text{IF } (lmin_t \text{ AND } \phi(ma_{t-1}^n(p_t)) \leq / \geq y) \text{ THEN } \dots \quad (6.1a)$$

$$\text{IF } (lmax_t \text{ AND } \phi(ma_{t-1}^n(p_t)) \leq / \geq z) \text{ THEN } \dots \quad (6.1b)$$

It is extremely difficult to develop intuition for these trading rules. By definition, there is an element of path dependence. However, for the price crossover rule, the resulting trading rule does not transform the trade distribution in any way. Rather, it samples it differently. For the trade reduction rule, it is feasible that the information uncovered by the resulting trading rule might vary according to the location of the buy/sell signals. This is because the location of the buy/sell signals captures something of the extent to which the price has become increasingly cheap or expensive relative to before. Since the trade reduction rule allows a trade to stay in the market for longer, this allows this information to persist for longer and so there is time to have an effect on the price. With reference to Figure 4.2 on page 55 for example, this means that there might be significant predictive information in the extremes of the trade distribution. More generally, the probability of ending up in the extremes of the trade distribution is lower the closer the trade entry point is to ± 1 . Thus, the probability of a trade being exited due to a loss is also lower. Because of this, the results for the trade reduction rule might be explained by a subset of trades that dominate the rule. While unlikely it cannot be discounted. This can be investigated by incrementally removing trades where the threshold for removing trades is stepped in, say, 0.1 increments. It is also desirable to have a sense of how the trade reduction rule sequences the buy/sell signals. One approach is to construct a decision tree and to relate the sequencing of the buy/sell signals to the probability of an event occurring. This makes it possible to define trading rules that, say, trade if and only if there is a 60% chance of a win.

As stated, moving average trading suffers from a stock selection problem. This is a difficult problem and one that limits the viability of moving average trading as a trading system. The proposed solution is to record $VP(k,i)$ for each stock at each time step such that each stock has a $VP(k,i)$ signature. Ranking each stock at each time step according to some function of

its signature is a potential method by which to solve the stock selection problem. A possible function is to weight the most recent $VP(k,i)$ values by a Gaussian or Epanechnikov kernel. However, implicit in this is the assumption that the maximally ranked stocks are those most suited to moving average trading. This needs to be investigated. It is also implicit that if the maximally ranked stocks are those most suited to moving average trading, they are the most likely to exhibit persistence in price direction. Consequently, the ranking scheme might be a proxy for price impact. This can be tested by ranking each stock based on a measure of price impact such as the Amihud Ratio (Amihud (2002)). Both ranking schemes need to be tested using the trade reduction and positive autocorrelation rules. This should provide insight into how to automate the trading rules with respect to stock selection.

It is also necessary to investigate the stock selection problem along with a trading strategy. It was concluded during the course of the thesis that the preferred strategy is to hold the capital value of each trade constant so that each trade earns the mean return per trade by definition. This breaks up serial dependence in the order of the trades and so losses are not compounded. Assuming the trading rules are profitable, the strategy is one where in the limit as the number of trades tends to infinity, $E[R_t] > 0$. The cash profit is then $kn\bar{R}$ where k is the capital value of each trade, n is the number of trades and \bar{R} is the mean return per trade. The advantage is that the strategy scales up and is massively parallel. Given unlimited capital, it is possible to trade the whole of the investment universe using multiple instances of the same trading rule parameterised in the moving average length. However, the problem is how to grow the scale of the strategy given limited capital. This is a non-trivial problem and needs to be addressed. Solving it is another step towards automating the trading rules.

As a test for market efficiency, the variable of interest is the breakeven costs. It is necessary to develop a test to test for the difference in the breakeven costs between trading rules. For a bootstrap test, this is equivalent to specifying the distribution under the null hypothesis of no difference. The breakeven cost is also the preferred way to test the statistical significance of the trading rules relative to the price crossover rule. It is also necessary to test the breakeven costs after adjusting for risk. Risk adjustment needs to be integrated into the breakeven costs test procedure.

A general issue that arises in response to the positive autocorrelation rule is whether positive autocorrelation dominates the price. One way to test this is to extend the trade reduction rule to include a test for positive autocorrelation and to test the difference between the resulting trading rule and the positive autocorrelation rule. This is equivalent to testing the difference between the price crossover and trade reduction rules where both trading rules include a test for positive autocorrelation. If positive autocorrelation dominates the price in that it induces consistently winning trades, the relationship between the price crossover and trade reduction rules of Chapter 4 might break down. If so, potentially at least, there is evidence that positive autocorrelation dominates the price. However, it is not clear what to expect in the event of a losing trade. If the price continues to trend, since the price crossover rule exits trades earlier than the trade reduction rule, the price crossover rule might perform better. The relationship of Chapter 4 might reverse. This should provide insight into the preferred choice of trading rule. There is also some similarity between this and testing for price impact.

A major contribution of this thesis is to establish a link between moving average trading and market microstructure. There is evidence that the positive autocorrelation rule is picking up asymmetric information. The positive autocorrelation rule is a method by which to peel back

the price and to identify events where there should be some change in the trading patterns in the high frequency data. It is thought that there should be evidence of adverse selection and price impact. This needs to be investigated. It might even be possible to recover the anatomy of a profitable trading opportunity at the microstructure level. There is also evidence that the price crossover and trade reduction rules are picking up changes in liquidity on the sell side when shorting the FTSE Small Cap and FTSE Fledgling portfolios. In general, naked shorts are the most difficult trades to execute consistently and successfully. Whether liquidity is a significant factor in determining the success of a short trade needs to be investigated. If it is, shorting might be more straightforward than it otherwise appears.

On a more speculative note, the structure of the positive autocorrelation rule is ideally suited to optimisation using genetic algorithms (Goldberg (1989), Mitchell (1998)). This can also be extended to include the trade reduction rule in the sense that the choice of trading rule is also a parameter. In the best case, the trading rule and stock selection problem can be treated as a joint optimisation problem that optimises the trading rule together with the optimal portfolio. Alternatively, optimising the trading rules on their own should provide some insight into the optimal parameters and hence into the type of trading rule most suited to testing for market efficiency. There may also be some value in combining the positive autocorrelation rule with time series forecasts (see, for example, Tsay (2010)). It is also worth noting that the bounded moving average reduces to a fuzzy Takagi-Sugeno model (Takagi and Sugeno (1985)). This may or may not provide a route in fuzzy modelling and associated technologies such as fuzzy neural networks.

Appendix

A.1 FTSE 100 Portfolio (Chapters 3 and 4)

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
3I Group	29-Mar-94	30-Jun-09		•	•	•
Admiral Group	22-Sep-04	30-Jun-09			•	•
Amec	22-Dec-82	30-Jun-09		•	•	•
Anglo American	31-Dec-90	30-Jun-09		•	•	•
Antofagasta	05-Jul-82	30-Jun-09		•	•	•
Associated British Foods	01-Jan-65	30-Jun-09	•	•	•	•
Astrazeneca	28-May-93	30-Jun-09		•	•	•
Autonomy Corporation	30-Oct-00	30-Jun-09			•	•
Aviva	01-Jan-65	30-Jun-09	•	•	•	•
BAE Systems	19-Feb-81	30-Jun-09		•	•	•
Balfour Beatty	01-Jan-65	30-Jun-09	•	•	•	•
Barclays	01-Jan-65	30-Jun-09	•	•	•	•
BG Group	05-Dec-86	30-Jun-09		•	•	•
BHP Billiton	25-Jul-97	30-Jun-09			•	•
BP	01-Jan-65	30-Jun-09	•	•	•	•
British Airways	10-Feb-87	30-Jun-09		•	•	•
British American Tobacco	01-Jan-65	30-Jun-09	•	•	•	•
British Land	01-Jan-65	30-Jun-09	•	•	•	•
British Sky Broadcasting	07-Dec-94	30-Jun-09		•	•	•
BT Group	30-Nov-84	30-Jun-09		•	•	•
Bunzl	01-Jan-65	30-Jun-09	•	•	•	•
Cable & Wireless	04-Nov-81	30-Jun-09		•	•	•
Cadbury	01-Jan-65	30-Jun-09	•	•	•	•
Cairn Energy	21-Dec-88	30-Jun-09		•	•	•
Capita Group	24-Apr-89	30-Jun-09		•	•	•
Carnival	20-Oct-00	30-Jun-09			•	•
Centrica	14-Feb-97	30-Jun-09			•	•
Cobham	01-Jan-65	30-Jun-09	•	•	•	•
Compass Group	01-Feb-01	30-Jun-09			•	•
Diageo	01-Jan-65	30-Jun-09	•	•	•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
Eurasian Natural Resources	06-Dec-07	30-Jun-09			•	•
Experian	06-Oct-06	30-Jun-09			•	•
Fresnillo	08-May-08	30-Jun-09			•	•
Friends Provident Group	06-Jul-01	30-Jun-09			•	•
G4S	10-Jun-96	30-Jun-09			•	•
Glaxosmithkline	01-Jan-65	30-Jun-09	•	•	•	•
Hammerson	01-Jan-65	30-Jun-09	•	•	•	•
Home Retail Group	01-Jan-65	30-Jun-09	•	•	•	•
HSBC Holdings	09-Jul-92	30-Jun-09		•	•	•
ICAP	16-Nov-98	30-Jun-09			•	•
Imperial Tobacco Group	30-Sep-96	30-Jun-09			•	•
Inmarsat	16-Jun-05	30-Jun-09			•	•
Intercontinental Hotels Group	28-Mar-03	30-Jun-09			•	•
International Power	11-Mar-91	30-Jun-09		•	•	•
Intertek Group	23-May-02	30-Jun-09			•	•
Invensys	29-Mar-72	30-Jun-09	•	•	•	•
Johnson Matthey	01-Jan-65	30-Jun-09	•	•	•	•
Kazakhmys	06-Oct-05	30-Jun-09			•	•
Kingfisher	24-Nov-82	30-Jun-09		•	•	•
Land Securities Group	01-Jan-65	30-Jun-09	•	•	•	•
Legal & General Group	01-Jan-65	30-Jun-09	•	•	•	•
Liberty International	29-Jul-92	30-Jun-09		•	•	•
Lloyds Banking Group	28-Dec-95	30-Jun-09			•	•
London Stock Exchange Group	21-Jul-00	30-Jun-09			•	•
Lonmin	01-Jan-65	30-Jun-09	•	•	•	•
Man Group	06-Oct-94	30-Jun-09		•	•	•
Marks & Spencer Group	01-Jan-65	30-Jun-09	•	•	•	•
Morrison (Wm) Supermarkets	01-Jan-69	30-Jun-09	•	•	•	•
National Grid	08-Dec-95	30-Jun-09			•	•
Next	01-Jan-65	30-Jun-09	•	•	•	•
Old Mutual	09-Jul-99	30-Jun-09			•	•
Pearson	20-Aug-69	30-Jun-09	•	•	•	•
Pennon Group	11-Dec-89	30-Jun-09		•	•	•
Petrofac	03-Oct-05	30-Jun-09			•	•
Prudential	01-Jan-65	30-Jun-09	•	•	•	•
Randgold Resources	30-Jun-97	30-Jun-09			•	•
Reckitt Benckiser Group	01-Jan-65	30-Jun-09	•	•	•	•
Reed Elsevier	01-Jan-65	30-Jun-09	•	•	•	•
Rexam	01-Jan-65	30-Jun-09	•	•	•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
Rio Tinto	01-Jan-65	30-Jun-09	•	•	•	•
Rolls-Royce Group	19-May-87	30-Jun-09		•	•	•
Royal Bank Of Scotland Group	01-Jan-65	30-Jun-09	•	•	•	•
Royal Dutch Shell B	01-Jan-65	30-Jun-09	•	•	•	•
Rsa Insurance Group	01-Jan-65	30-Jun-09	•	•	•	•
Sabmiller	26-Feb-99	30-Jun-09			•	•
Sage Group	13-Dec-89	30-Jun-09		•	•	•
Sainsbury (J)	18-Jul-73	30-Jun-09	•	•	•	•
Schroders	01-Jan-65	30-Jun-09	•	•	•	•
Schroders N-V	08-May-86	30-Jun-09		•	•	•
Scottish & Southern Energy	17-Jun-91	30-Jun-09		•	•	•
Serco Group	11-May-88	30-Jun-09		•	•	•
Severn Trent	11-Dec-89	30-Jun-09		•	•	•
Shire	14-Feb-96	30-Jun-09			•	•
Smith & Nephew	01-Jan-65	30-Jun-09	•	•	•	•
Smiths Group	01-Jan-65	30-Jun-09	•	•	•	•
Standard Chartered	01-Jan-65	30-Jun-09	•	•	•	•
Standard Life	07-Jul-06	30-Jun-09			•	•
Tesco	01-Jan-65	30-Jun-09	•	•	•	•
Thomas Cook Group	30-Dec-04	30-Jun-09			•	•
Thomson Reuters	01-Jun-84	30-Jun-09		•	•	•
Tui Travel	06-Jan-82	30-Jun-09		•	•	•
Tullow Oil	04-Oct-89	30-Jun-09		•	•	•
Unilever	01-Jan-65	30-Jun-09	•	•	•	•
United Utilities Group	11-Dec-89	30-Jun-09		•	•	•
Vedanta Resources	04-Dec-03	30-Jun-09			•	•
Vodafone Group	25-Oct-88	30-Jun-09		•	•	•
Wolseley	01-Jan-65	30-Jun-09	•	•	•	•
Wpp	14-Apr-71	30-Jun-09	•	•	•	•
Xstrata	19-Mar-02	30-Jun-09			•	•

Period 1 refers to the test period 01-Jan-1965 to 31-Dec-1979. Period 2 refers to the test period 01-Jan-1980 to 31-Dec-1994. Period 3 refers to the test period 01-Jan-1995 to 30-Jun-2009. Period All refers to the test period 01-Jan-1965 to 30-Jun-2009. A • symbol indicates inclusion in the test period.

A.2 FTSE 250 Portfolio (Chapters 3 and 4)

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
3I Infrastructure	26-Feb-07	30-Jun-09			•	•
888 Holdings	28-Sep-05	30-Jun-09			•	•
Aberdeen Asset Management	27-Mar-91	30-Jun-09		•	•	•
Aegis Group	21-Jan-83	30-Jun-09		•	•	•
Aggreko	26-Sep-97	30-Jun-09			•	•
Amlin	25-Nov-93	30-Jun-09		•	•	•
Aquarius Platinum	04-Oct-99	30-Jun-09			•	•
Arm Holdings	23-Apr-98	30-Jun-09			•	•
Arriva	01-Jan-69	30-Jun-09	•	•	•	•
Ashmore Group	11-Oct-06	30-Jun-09			•	•
Ashtead Group	04-Dec-86	30-Jun-09		•	•	•
Atkins (WS)	24-Jul-96	30-Jun-09			•	•
Aveva Group	04-Dec-96	30-Jun-09			•	•
Babcock International Group	11-Aug-89	30-Jun-09		•	•	•
Barr (AG)	27-Aug-69	30-Jun-09	•	•	•	•
Barratt Developments	01-Jan-69	30-Jun-09	•	•	•	•
BBA Aviation	01-Jan-65	30-Jun-09	•	•	•	•
Beazley	11-Nov-02	30-Jun-09			•	•
Bellway	16-May-79	30-Jun-09	•	•	•	•
Berkeley Group Holdings	18-Jul-84	30-Jun-09		•	•	•
Big Yellow Group	05-May-00	30-Jun-09			•	•
Bluebay Asset Management	16-Nov-06	30-Jun-09			•	•
Bodycote	12-Jan-72	30-Jun-09	•	•	•	•
Bovis Homes Group	08-Dec-97	30-Jun-09			•	•
Brewin Dolphin Holdings	08-Jun-94	30-Jun-09		•	•	•
Brit Insurance Holdings	27-Oct-95	30-Jun-09			•	•
Britvic	08-Dec-05	30-Jun-09			•	•
Brown (N) Group	01-Apr-70	30-Jun-09	•	•	•	•
BSS Group	01-Jan-65	30-Jun-09	•	•	•	•
BTG	05-Jul-95	30-Jun-09			•	•
Burberry Group	11-Jul-02	30-Jun-09			•	•
Carillion	29-Jul-99	30-Jun-09			•	•
Carpetright	22-Jun-93	30-Jun-09		•	•	•
Carphone Warehouse Group	13-Jul-00	30-Jun-09			•	•
Catlin Group	31-Mar-04	30-Jun-09			•	•
Chaucer Holdings	23-Nov-93	30-Jun-09		•	•	•
Chemring Group	01-Jan-69	30-Jun-09	•	•	•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
Chloride Group	01-Jan-65	30-Jun-09	•	•	•	•
Close Brothers Group	01-Jan-65	30-Jun-09	•	•	•	•
Computacenter	20-May-98	30-Jun-09			•	•
Connaught	27-Nov-98	30-Jun-09			•	•
Cookson Group	01-Jan-65	30-Jun-09	•	•	•	•
Cranswick	04-Dec-85	30-Jun-09		•	•	•
Croda International	01-Jan-65	30-Jun-09	•	•	•	•
CSR	25-Feb-04	30-Jun-09			•	•
Daejan Holdings	01-Jan-65	30-Jun-09	•	•	•	•
Daily Mail & General Trust	01-Jan-65	30-Jun-09	•	•	•	•
Dairy Crest Group	27-Aug-96	30-Jun-09			•	•
Dana Petroleum	05-Jun-95	30-Jun-09			•	•
Davis Service Group	01-Jan-65	30-Jun-09	•	•	•	•
De La Rue	01-Jan-65	30-Jun-09	•	•	•	•
Debenhams	03-May-06	30-Jun-09			•	•
Dechra Pharmaceuticals	20-Sep-00	30-Jun-09			•	•
Derwent London	10-Aug-84	30-Jun-09		•	•	•
Dignity	01-Apr-04	30-Jun-09			•	•
Dimension Data Holdings	18-Jul-00	30-Jun-09			•	•
Domino Printing Sciences	01-May-85	30-Jun-09		•	•	•
Dominos Pizza	23-Nov-99	30-Jun-09			•	•
Drax Group	14-Dec-05	30-Jun-09			•	•
DSG International	01-Jan-65	30-Jun-09	•	•	•	•
Dunelm Group	18-Oct-06	30-Jun-09			•	•
Eaga	06-Jun-07	30-Jun-09			•	•
Easyjet	14-Nov-00	30-Jun-09			•	•
Electrocomponents	05-Jul-67	30-Jun-09	•	•	•	•
Emerald Energy	15-Nov-93	30-Jun-09		•	•	•
Enterprise Inns	03-Nov-95	30-Jun-09			•	•
Euromoney Institutional Investors	23-Jun-86	30-Jun-09		•	•	•
Evolution Group	25-Jun-97	30-Jun-09			•	•
F&C Asset Management	02-Sep-83	30-Jun-09		•	•	•
Ferrexpo	14-Jun-07	30-Jun-09			•	•
Fidessa Group	06-Jun-97	30-Jun-09			•	•
Filtrona	03-Jun-05	30-Jun-09			•	•
First Group	15-Jun-95	30-Jun-09			•	•
Fisher (James) & Sons	01-Jan-65	30-Jun-09	•	•	•	•
Forth Ports	20-Mar-92	30-Jun-09		•	•	•
Galiform	16-Jul-92	30-Jun-09		•	•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
Game Group	05-Jul-85	30-Jun-09		•	•	•
Genus	05-Jul-00	30-Jun-09			•	•
GKN	01-Jan-65	30-Jun-09	•	•	•	•
Go-Ahead Group	06-May-94	30-Jun-09		•	•	•
Great Portland Estates	01-Jan-65	30-Jun-09	•	•	•	•
Greene King	01-Jan-65	30-Jun-09	•	•	•	•
Greggs	03-May-84	30-Jun-09		•	•	•
Halfords Group	02-Jun-04	30-Jun-09			•	•
Halma	19-Jan-72	30-Jun-09	•	•	•	•
Hargreaves Lansdown	14-May-07	30-Jun-09			•	•
Hays	25-Oct-89	30-Jun-09		•	•	•
Helical Bar	01-Jan-69	30-Jun-09	•	•	•	•
Henderson Group	22-Dec-03	30-Jun-09			•	•
Hikma Pharmaceuticals	31-Oct-05	30-Jun-09			•	•
Hiscox	30-Jun-95	30-Jun-09			•	•
HMV Group	08-May-02	30-Jun-09			•	•
Hochschild Mining	02-Nov-06	30-Jun-09			•	•
Homeserve	26-Nov-91	30-Jun-09		•	•	•
Hunting	29-Jul-70	30-Jun-09	•	•	•	•
IG Group Holdings	27-Apr-05	30-Jun-09			•	•
IMI	06-Apr-66	30-Jun-09	•	•	•	•
Inchcape	01-Jan-65	30-Jun-09	•	•	•	•
Informa	16-Apr-98	30-Jun-09			•	•
Intermediate Capital Group	31-May-94	30-Jun-09		•	•	•
Interserve	01-Jan-65	30-Jun-09	•	•	•	•
Investec	19-Jul-02	30-Jun-09			•	•
ITV	01-Jan-65	30-Jun-09	•	•	•	•
Jardine Lloyd Thompson Group	16-Oct-87	30-Jun-09		•	•	•
JKX Oil & Gas	11-Jul-95	30-Jun-09			•	•
Keller	04-May-94	30-Jun-09		•	•	•
Kesa Electricals	04-Jul-03	30-Jun-09			•	•
Kier Group	11-Dec-96	30-Jun-09			•	•
Ladbroke's	04-Oct-67	30-Jun-09	•	•	•	•
Lancashire Holdings	12-Dec-05	30-Jun-09			•	•
Logica	02-Nov-83	30-Jun-09		•	•	•
Marstons	01-Jan-65	30-Jun-09	•	•	•	•
McBride	06-Jul-95	30-Jun-09			•	•
Meggitt	01-Jan-69	30-Jun-09	•	•	•	•
Melrose	27-Oct-03	30-Jun-09			•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
Melrose Resources	17-Dec-99	30-Jun-09			•	•
Michael Page International	27-Mar-01	30-Jun-09			•	•
Micro Focus International	11-May-05	30-Jun-09			•	•
Millennium & Copthorne Hotels	24-Apr-96	30-Jun-09			•	•
Misys	11-Mar-87	30-Jun-09		•	•	•
Mitchells & Butlers	28-Mar-03	30-Jun-09			•	•
Mitie Group	01-Jan-69	30-Jun-09	•	•	•	•
Mondi	29-Jun-07	30-Jun-09			•	•
Moneysupermarket Dot Com	25-Jul-07	30-Jun-09			•	•
Morgan Crucible	01-Jan-65	30-Jun-09	•	•	•	•
Morgan Sindall	01-Jan-69	30-Jun-09	•	•	•	•
Mothercare	07-Jan-86	30-Jun-09		•	•	•
Mouchel Group	27-Jun-02	30-Jun-09			•	•
National Express Group	09-Dec-92	30-Jun-09		•	•	•
Northern Foods	01-Jan-65	30-Jun-09	•	•	•	•
Northumbrian Water Group	22-May-03	30-Jun-09			•	•
Novae Group	24-Nov-93	30-Jun-09		•	•	•
Pace	26-Jun-96	30-Jun-09			•	•
Partygaming	24-Jun-05	30-Jun-09			•	•
Paypoint	20-Sep-04	30-Jun-09			•	•
Persimmon	26-Apr-85	30-Jun-09		•	•	•
Peter Hambro Mining	26-Apr-02	30-Jun-09			•	•
Premier Farnell	01-Jun-66	30-Jun-09	•	•	•	•
Premier Foods	19-Jul-04	30-Jun-09			•	•
Premier Oil	21-Feb-73	30-Jun-09	•	•	•	•
Provident Financial	01-Jan-65	30-Jun-09	•	•	•	•
Punch Taverns	21-May-02	30-Jun-09			•	•
PV Crystalox Solar	05-Jun-07	30-Jun-09			•	•
PZ Cussons	01-Jan-65	30-Jun-09	•	•	•	•
Qinetiq Group	09-Feb-06	30-Jun-09			•	•
Rank Group	01-Jan-65	30-Jun-09	•	•	•	•
Rathbone Brothers	24-Sep-84	30-Jun-09		•	•	•
Redrow	16-May-94	30-Jun-09		•	•	•
Regus	16-Oct-00	30-Jun-09			•	•
Renishaw	02-Jun-83	30-Jun-09		•	•	•
Rentokil Initial	19-Mar-69	30-Jun-09	•	•	•	•
Restaurant Group	01-Jan-65	30-Jun-09	•	•	•	•
Rightmove	09-Mar-06	30-Jun-09			•	•
Robert Wiseman Dairies	25-Mar-94	30-Jun-09		•	•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
Rotork	24-Jul-68	30-Jun-09	•	•	•	•
RPS Group	28-Jul-87	30-Jun-09		•	•	•
Salamander Energy	29-Nov-06	30-Jun-09			•	•
Savills	20-Jul-88	30-Jun-09		•	•	•
SDL	06-Dec-99	30-Jun-09			•	•
Segro	01-Jan-65	30-Jun-09	•	•	•	•
Shaftesbury	19-Oct-87	30-Jun-09		•	•	•
Shanks Group	26-Feb-88	30-Jun-09		•	•	•
SIG	17-May-89	30-Jun-09		•	•	•
Smith (DS)	01-Jan-69	30-Jun-09	•	•	•	•
Soco International	28-May-97	30-Jun-09			•	•
Spectris	28-Nov-88	30-Jun-09		•	•	•
Spirax-Sarco Engineering	01-Jan-65	30-Jun-09	•	•	•	•
Spirent Communications	01-Jan-65	30-Jun-09	•	•	•	•
Sports Direct International	26-Feb-07	30-Jun-09			•	•
SSL International	13-Jul-90	30-Jun-09		•	•	•
St James Place	23-Aug-96	30-Jun-09			•	•
Stagecoach Group	26-Apr-93	30-Jun-09		•	•	•
Sthree	10-Nov-05	30-Jun-09			•	•
Stobart Group	26-Feb-02	30-Jun-09			•	•
Synergy Health	17-Aug-01	30-Jun-09			•	•
Tate & Lyle	01-Jan-65	30-Jun-09	•	•	•	•
Taylor Wimpey	01-Jan-65	30-Jun-09	•	•	•	•
Telecity Group	23-Oct-07	30-Jun-09			•	•
Tomkins	01-Jan-69	30-Jun-09	•	•	•	•
Travis Perkins	18-Sep-86	30-Jun-09		•	•	•
Tullett Prebon	13-Dec-06	30-Jun-09			•	•
UK Commercial Property Trust	29-Aug-06	30-Jun-09			•	•
Ultra Electronics Holdings	02-Oct-96	30-Jun-09			•	•
United Business Media	01-Jan-65	30-Jun-09	•	•	•	•
Vectura Group	01-Jul-04	30-Jun-09			•	•
Venture Production	18-Mar-02	30-Jun-09			•	•
Victrex	20-Dec-95	30-Jun-09			•	•
VT Group	16-Mar-88	30-Jun-09		•	•	•
Weir Group	01-Jan-65	30-Jun-09	•	•	•	•
Wellstream Holdings	25-Apr-07	30-Jun-09			•	•
Wetherspoon (JD)	29-Oct-92	30-Jun-09		•	•	•
WH Smith	01-Jan-65	30-Jun-09	•	•	•	•
Whitbread	01-Jan-65	30-Jun-09	•	•	•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
William Hill	14-Jun-02	30-Jun-09			•	•
Wood Group (John)	28-May-02	30-Jun-09			•	•
Xchanging	24-Apr-07	30-Jun-09			•	•
Yell Group	09-Jul-03	30-Jun-09			•	•

Period 1 refers to the test period 01-Jan-1965 to 31-Dec-1979. Period 2 refers to the test period 01-Jan-1980 to 31-Dec-1994. Period 3 refers to the test period 01-Jan-1995 to 30-Jun-2009. Period All refers to the test period 01-Jan-1965 to 30-Jun-2009. A • symbol indicates inclusion in the test period.

A.3 FTSE Small Cap Portfolio (Chapters 3 and 4)

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
AGA Rangemaster Group	01-Jan-65	30-Jun-09	•	•	•	•
Air Partner	03-Nov-89	30-Jun-09		•	•	•
Alpha Pyrenees	28-Nov-05	30-Jun-09			•	•
Alphameric	03-Aug-84	30-Jun-09		•	•	•
Alterian	19-Jul-00	30-Jun-09			•	•
Anglo Eastern Plantations	17-May-85	30-Jun-09		•	•	•
Anglo Pacific Group	29-Mar-84	30-Jun-09		•	•	•
Anite	27-Jun-73	30-Jun-09	•	•	•	•
Antisoma	15-Dec-99	30-Jun-09			•	•
Arena Leisure	15-Nov-72	30-Jun-09	•	•	•	•
Ark Therapeutics Group	02-Mar-04	30-Jun-09			•	•
Ashley (Laura) Holdings	04-Dec-85	30-Jun-09		•	•	•
Assura Group	29-Oct-03	30-Jun-09			•	•
Avis Europe	03-Apr-97	30-Jun-09			•	•
Axis-Shield	22-Sep-93	30-Jun-09		•	•	•
Bloomsbury Publishing	22-Jun-94	30-Jun-09		•	•	•
Braemar Shipping Services	26-Nov-97	30-Jun-09			•	•
Brammer (H)	01-Jan-65	30-Jun-09	•	•	•	•
British Polythene Industries	07-Apr-65	30-Jun-09	•	•	•	•
Brixton	01-Jan-65	30-Jun-09	•	•	•	•
Business Post Group	02-Jul-93	30-Jun-09		•	•	•
Camellia	01-Jan-69	30-Jun-09	•	•	•	•
Capital & Regional	12-Dec-86	30-Jun-09		•	•	•
Care UK	16-Jul-86	30-Jun-09		•	•	•
Castings	01-Jan-69	30-Jun-09	•	•	•	•
Centaur Media	09-Mar-04	30-Jun-09			•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
Charles Stanley Group	01-Jan-69	30-Jun-09	•	•	•	•
Charles Taylor Consulting	09-Oct-96	30-Jun-09			•	•
Chesnara	19-May-04	30-Jun-09			•	•
Chime Communications	05-Jan-90	30-Jun-09		•	•	•
Chrysalis	05-Mar-69	30-Jun-09	•	•	•	•
Cineworld Group	26-Apr-07	30-Jun-09			•	•
Clarke (T)	01-Jan-69	30-Jun-09	•	•	•	•
Clarkson	27-Jun-86	30-Jun-09		•	•	•
CLS Holdings	26-May-94	30-Jun-09		•	•	•
Collins Stewart	23-Oct-00	30-Jun-09			•	•
Communis	24-Jun-94	30-Jun-09		•	•	•
Consort Medical	24-Nov-82	30-Jun-09		•	•	•
Costain Group	01-Jan-65	30-Jun-09	•	•	•	•
Delta	01-Jan-65	30-Jun-09	•	•	•	•
Development Securities	07-May-85	30-Jun-09		•	•	•
Devro	29-Jun-93	30-Jun-09		•	•	•
Diploma	01-Jan-69	30-Jun-09	•	•	•	•
DTZ Holdings	29-Jul-87	30-Jun-09		•	•	•
E2V Technologies	19-Jul-04	30-Jun-09			•	•
Elementis	01-Jan-65	30-Jun-09	•	•	•	•
Emblaze	17-Oct-96	30-Jun-09			•	•
Fenner	01-Jan-65	30-Jun-09	•	•	•	•
Fiberweb	16-Nov-06	30-Jun-09			•	•
Findel	01-Jan-65	30-Jun-09	•	•	•	•
Fortune Oil	27-Sep-89	30-Jun-09		•	•	•
French Connection Group	07-Nov-83	30-Jun-09		•	•	•
Fuller Smith & Turner	05-Nov-80	30-Jun-09		•	•	•
Future	17-Jun-99	30-Jun-09			•	•
Galliford Try	01-Jan-69	30-Jun-09	•	•	•	•
Gem Diamonds	13-Feb-07	30-Jun-09			•	•
Gleeson (MJ) Group	01-Jan-65	30-Jun-09	•	•	•	•
Goldenport Holdings	31-Mar-06	30-Jun-09			•	•
Goldshield Group	11-Jun-98	30-Jun-09			•	•
Goodwin	01-Jan-69	30-Jun-09	•	•	•	•
Grainger	02-Mar-83	30-Jun-09		•	•	•
Hampson Industries	01-Jan-69	30-Jun-09	•	•	•	•
Hardy Oil & Gas	06-Jun-05	30-Jun-09			•	•
Hardy Underwriting Bermuda	27-Dec-96	30-Jun-09			•	•
Headlam Group	01-Jan-69	30-Jun-09	•	•	•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
Helphire Group	21-Mar-97	30-Jun-09			•	•
Hill & Smith Holdings	02-Apr-69	30-Jun-09	•	•	•	•
Hilton Food Group	16-May-07	30-Jun-09			•	•
Hogg Robinson Group	06-Oct-06	30-Jun-09			•	•
Holidaybreak	09-Jul-91	30-Jun-09		•	•	•
Hornby	17-Dec-86	30-Jun-09		•	•	•
Hyder Consulting	09-Aug-72	30-Jun-09	•	•	•	•
Imagination Technologies Group	05-Jul-94	30-Jun-09		•	•	•
Innovation Group	01-Jun-00	30-Jun-09			•	•
Intec Telecom Systems	09-Jun-00	30-Jun-09			•	•
International Ferro Metals	29-Sep-05	30-Jun-09			•	•
International Personal Finance	13-Jul-07	30-Jun-09			•	•
IP Group	14-Oct-03	30-Jun-09			•	•
IRP Property Investments	28-Apr-04	30-Jun-09			•	•
Isis Property Trust	26-Sep-03	30-Jun-09			•	•
ITE Group	10-Jan-94	30-Jun-09		•	•	•
JD Sports Fashion	21-Oct-96	30-Jun-09			•	•
JJB Sports	17-Nov-94	30-Jun-09		•	•	•
Johnston Press	28-Apr-88	30-Jun-09		•	•	•
Kcom Group	09-Jul-99	30-Jun-09			•	•
Kewill	16-Sep-85	30-Jun-09		•	•	•
Kofax	02-Apr-96	30-Jun-09			•	•
Laird	01-Jan-65	30-Jun-09	•	•	•	•
Lamprell	10-Oct-06	30-Jun-09			•	•
Lavendon Group	09-Oct-96	30-Jun-09			•	•
Lookers	27-Jun-73	30-Jun-09	•	•	•	•
Low & Bonar	01-Jan-65	30-Jun-09	•	•	•	•
LSL Property Services	15-Nov-06	30-Jun-09			•	•
Luminar Group Holdings	17-May-96	30-Jun-09			•	•
Management Consulting Group	16-Feb-87	30-Jun-09		•	•	•
Marshalls	01-Jan-69	30-Jun-09	•	•	•	•
McKay Securities	01-Jan-69	30-Jun-09	•	•	•	•
Mears Group	03-Oct-96	30-Jun-09			•	•
Mecom Group	22-Mar-05	30-Jun-09			•	•
Medicx Fund	27-Oct-06	30-Jun-09			•	•
Menzies (John)	01-Jan-65	30-Jun-09	•	•	•	•
Mucklow (A & J) Group	01-Jan-65	30-Jun-09	•	•	•	•
MWB Group Holdings	02-Apr-08	30-Jun-09			•	•
NCC Group	08-Jul-04	30-Jun-09			•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
Northgate	01-Jan-69	30-Jun-09	•	•	•	•
Optos	09-Feb-06	30-Jun-09			•	•
Oxford Biomedica	12-Dec-96	30-Jun-09			•	•
Oxford Instruments	18-Oct-83	30-Jun-09		•	•	•
Paragon Group Of Companies	25-Sep-85	30-Jun-09		•	•	•
Pendragon	10-Nov-89	30-Jun-09		•	•	•
Phoenix IT Group	10-Nov-04	30-Jun-09			•	•
Photo-Me International	01-Jan-65	30-Jun-09	•	•	•	•
Primary Health Properties	18-Mar-96	30-Jun-09			•	•
Prostrakan Group	13-Jun-05	30-Jun-09			•	•
Psion	11-Mar-88	30-Jun-09		•	•	•
Quintain Estates & Development	22-Jul-96	30-Jun-09			•	•
REA Holdings	01-Jan-69	30-Jun-09	•	•	•	•
Real Estate Opportunities	29-May-01	30-Jun-09			•	•
Renovo Group	06-Apr-06	30-Jun-09			•	•
Rensburg Sheppards	06-Apr-88	30-Jun-09		•	•	•
Ricardo	01-Jan-69	30-Jun-09	•	•	•	•
RM	13-Dec-94	30-Jun-09		•	•	•
Robert Walters	05-Jul-00	30-Jun-09			•	•
Rok	19-Aug-81	30-Jun-09		•	•	•
RPC Group	27-May-93	30-Jun-09		•	•	•
Safestore Holdings	08-Mar-07	30-Jun-09			•	•
Scott Wilson Group	14-Mar-06	30-Jun-09			•	•
Senior	01-Jan-65	30-Jun-09	•	•	•	•
Sepura	30-Jul-07	30-Jun-09			•	•
Severfield-Rowen	01-Jul-88	30-Jun-09		•	•	•
Smiths News	29-Aug-06	30-Jun-09			•	•
Southern Cross Healthcare	06-Jul-06	30-Jun-09			•	•
Speedy Hire	21-Jun-89	30-Jun-09		•	•	•
Spice	25-Aug-04	30-Jun-09			•	•
Sportech	28-Feb-86	30-Jun-09		•	•	•
Spring Group	01-Jan-69	30-Jun-09	•	•	•	•
St Ives	02-Oct-85	30-Jun-09		•	•	•
St Modwen Properties	25-Apr-86	30-Jun-09		•	•	•
STV Group	01-Jan-69	30-Jun-09	•	•	•	•
Ted Baker	23-Jul-97	30-Jun-09			•	•
Telecom Plus	16-Oct-98	30-Jun-09			•	•
Thorntons	23-May-88	30-Jun-09		•	•	•
Topps Tiles	30-May-97	30-Jun-09			•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
Town Centre Securities	01-Jan-69	30-Jun-09	•	•	•	•
Tribal Group	22-Feb-01	30-Jun-09			•	•
Trinity Mirror	01-Jan-65	30-Jun-09	•	•	•	•
TT Electronics	01-Jan-69	30-Jun-09	•	•	•	•
UK Coal	04-Jun-93	30-Jun-09		•	•	•
Umeco	07-Jul-89	30-Jun-09		•	•	•
Unite Group	31-May-99	30-Jun-09			•	•
UTV Media	19-Dec-86	30-Jun-09		•	•	•
Vitec Group	13-Dec-72	30-Jun-09	•	•	•	•
VP	11-Apr-73	30-Jun-09	•	•	•	•
Wilmington Group	05-Dec-95	30-Jun-09			•	•
Wincanton	17-May-01	30-Jun-09			•	•
Wolfson Microelectronics	15-Oct-03	30-Jun-09			•	•
Workspace Group	14-Dec-93	30-Jun-09		•	•	•
WSP Group	28-Sep-87	30-Jun-09		•	•	•
Yule Catto & Co	20-Oct-71	30-Jun-09	•	•	•	•

Period 1 refers to the test period 01-Jan-1965 to 31-Dec-1979. Period 2 refers to the test period 01-Jan-1980 to 31-Dec-1994. Period 3 refers to the test period 01-Jan-1995 to 30-Jun-2009. Period All refers to the test period 01-Jan-1965 to 30-Jun-2009. A • symbol indicates inclusion in the test period.

A.4 FTSE Fledgling Portfolio (Chapters 3 and 4)

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
4Imprint Group	01-Jan-65	30-Jun-09	•	•	•	•
600 Group	01-Jan-65	30-Jun-09	•	•	•	•
Abbeycrest	21-May-85	30-Jun-09		•	•	•
Acal	15-Jun-88	30-Jun-09		•	•	•
Accident Exchange Group	01-May-02	30-Jun-09			•	•
Advantage Property Income Trust	07-Feb-05	30-Jun-09			•	•
AEA Technology	25-Sep-96	30-Jun-09			•	•
Alexandra	30-Jan-85	30-Jun-09		•	•	•
Alexon Group	01-Jan-65	30-Jun-09	•	•	•	•
Alumasc Group	29-May-86	30-Jun-09		•	•	•
Anglesey Mining	03-Aug-88	30-Jun-09		•	•	•
Arc International	20-Sep-00	30-Jun-09			•	•
Associated British Engineering	01-Jan-65	30-Jun-09	•	•	•	•
Asterand	28-Jul-00	30-Jun-09			•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
Avon Rubber	01-Jan-65	30-Jun-09	•	•	•	•
Axa Property Trust	29-Mar-05	30-Jun-09			•	•
Beale	10-Mar-95	30-Jun-09			•	•
Berkeley Technology	08-Jan-85	30-Jun-09		•	•	•
Bisichi Mining	01-Jan-69	30-Jun-09	•	•	•	•
Blacks Leisure Group	28-Oct-70	30-Jun-09	•	•	•	•
Caffyns	01-Jan-65	30-Jun-09	•	•	•	•
Carclo	01-Jan-69	30-Jun-09	•	•	•	•
Cardiff Property	01-Jan-69	30-Jun-09	•	•	•	•
Carrs Milling Industries	24-May-72	30-Jun-09	•	•	•	•
Celsis International	05-Jul-93	30-Jun-09		•	•	•
City Of London Group	01-Jul-88	30-Jun-09		•	•	•
Clinton Cards	04-May-88	30-Jun-09		•	•	•
CML Microsystems	07-Feb-84	30-Jun-09		•	•	•
Coral Products	12-Apr-95	30-Jun-09			•	•
Corin Group	08-May-02	30-Jun-09			•	•
Cosalt	23-Jun-71	30-Jun-09	•	•	•	•
Creightons	05-Sep-86	30-Jun-09		•	•	•
Creston	01-Jan-69	30-Jun-09	•	•	•	•
Dawson Holdings	16-Jun-95	30-Jun-09			•	•
Dee Valley Group	16-Dec-94	30-Jun-09		•	•	•
Dialight	08-Nov-93	30-Jun-09		•	•	•
DRS Data & Research Services	04-May-94	30-Jun-09		•	•	•
Dyson Group	01-Jan-65	30-Jun-09	•	•	•	•
Electronic Data Processing	27-Sep-85	30-Jun-09		•	•	•
Filtronic	21-Oct-94	30-Jun-09		•	•	•
GB Group	28-May-93	30-Jun-09		•	•	•
Gresham Computing	28-Jun-84	30-Jun-09		•	•	•
Harvard International	30-Sep-87	30-Jun-09		•	•	•
Harvey Nash Group	02-Apr-97	30-Jun-09			•	•
Havelock Europa	30-Mar-84	30-Jun-09		•	•	•
Haynes Publishing Group	05-Dec-79	30-Jun-09	•	•	•	•
Heywood Williams Group	01-Jan-69	30-Jun-09	•	•	•	•
Highcroft Investments	17-Jun-70	30-Jun-09	•	•	•	•
Highway Capital	17-Mar-95	30-Jun-09			•	•
HR Owen	06-Feb-89	30-Jun-09		•	•	•
Jarvis	01-Jan-69	30-Jun-09	•	•	•	•
Jessops	28-Oct-04	30-Jun-09			•	•
Litho Supplies	19-Nov-93	30-Jun-09		•	•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
Local Shopping REIT (The)	26-Apr-07	30-Jun-09			•	•
London & Associated Properties	01-Jan-69	30-Jun-09	•	•	•	•
London Finance & Investment	01-Jan-69	30-Jun-09	•	•	•	•
Macfarlane Group	20-Jun-73	30-Jun-09	•	•	•	•
Mallet	19-Mar-87	30-Jun-09		•	•	•
Manganese Bronze Holdings	01-Jan-65	30-Jun-09	•	•	•	•
Microgen	14-Jan-83	30-Jun-09		•	•	•
Minerva	27-Nov-96	30-Jun-09			•	•
Molins	28-Jul-76	30-Jun-09	•	•	•	•
Morse	19-Mar-99	30-Jun-09			•	•
Moss Bros Group	01-Jan-65	30-Jun-09	•	•	•	•
MS International	01-Jan-69	30-Jun-09	•	•	•	•
Narborough Plantations	01-Jan-69	30-Jun-09	•	•	•	•
Nestor Healthcare Group	02-Dec-87	30-Jun-09		•	•	•
Network Technology	29-Jul-96	30-Jun-09			•	•
Norcros	13-Jul-07	30-Jun-09			•	•
North Midland Construction	01-Jan-69	30-Jun-09	•	•	•	•
Northamber	08-Jun-84	30-Jun-09		•	•	•
NXT	01-Jan-69	30-Jun-09	•	•	•	•
Office2Office	28-Jun-04	30-Jun-09			•	•
Parity Group	29-Jun-87	30-Jun-09		•	•	•
Parkwood Holdings	09-Dec-96	30-Jun-09			•	•
Phytopharm	24-Apr-96	30-Jun-09			•	•
Pochins	01-Jan-69	30-Jun-09	•	•	•	•
Porvair	04-May-88	30-Jun-09		•	•	•
Puricore	29-Jun-06	30-Jun-09			•	•
Queens Walk Investment	07-Dec-05	30-Jun-09			•	•
Raymarine	03-Dec-04	30-Jun-09			•	•
Renold	01-Jan-65	30-Jun-09	•	•	•	•
Ross Group	01-Jan-69	30-Jun-09	•	•	•	•
Rugby Estates Investment Trust	14-May-07	30-Jun-09			•	•
S & U	01-Jan-65	30-Jun-09	•	•	•	•
Sinclair Pharma	10-Dec-03	30-Jun-09			•	•
Skyepharm	23-Oct-87	30-Jun-09		•	•	•
Smart (J) & Co	01-Jan-69	30-Jun-09	•	•	•	•
Source Bioscience	01-Jan-69	30-Jun-09	•	•	•	•
Stanelco	14-Apr-88	30-Jun-09		•	•	•
Styles & Wood Group	06-Nov-06	30-Jun-09			•	•
Superglass Holdings	11-Jul-07	30-Jun-09			•	•

Stock	Start Date	End Date	Period 1	Period 2	Period 3	Period All
Tex Holdings	01-Dec-71	30-Jun-09	•	•	•	•
Titon Holdings	01-Feb-88	30-Jun-09		•	•	•
Torotrak	24-Jul-98	30-Jun-09			•	•
Total Systems	30-Mar-88	30-Jun-09		•	•	•
Trafficmaster	30-Mar-94	30-Jun-09		•	•	•
Treatt	21-Jun-89	30-Jun-09		•	•	•
Triad Group	20-Mar-96	30-Jun-09			•	•
Trifast	15-Feb-94	30-Jun-09		•	•	•
Uniq	01-Jan-65	30-Jun-09	•	•	•	•
Vernalis	30-Jun-92	30-Jun-09		•	•	•
Victoria	02-Oct-68	30-Jun-09	•	•	•	•
Vislink	18-Apr-79	30-Jun-09	•	•	•	•
Volex Group	01-Jan-65	30-Jun-09	•	•	•	•
Walker Crips Group	21-Aug-96	30-Jun-09			•	•
Warner Estate Holdings	01-Jan-69	30-Jun-09	•	•	•	•
Waterman Group	23-May-88	30-Jun-09		•	•	•
White Young Green	28-Apr-86	30-Jun-09		•	•	•
Worthington Group	01-Jan-69	30-Jun-09	•	•	•	•
Xaar	16-Oct-97	30-Jun-09			•	•
Zotefoams	27-Feb-95	30-Jun-09			•	•

Period 1 refers to the test period 01-Jan-1965 to 31-Dec-1979. Period 2 refers to the test period 01-Jan-1980 to 31-Dec-1994. Period 3 refers to the test period 01-Jan-1995 to 30-Jun-2009. Period All refers to the test period 01-Jan-1965 to 30-Jun-2009. A • symbol indicates inclusion in the test period.

A.5 Sharescope

Ionic Information Ltd

Ionic House

3 Bath Place, Rivington Street

London EC2A 3DR

<http://www.sharescope.co.uk>

A.6 Udata Technical Analyst

Udata

Hutchison House

5 Hester Road

London SW11 4AN

<http://www.udata.co.uk>

A.7 FTSE 100 Portfolio (Chapter 5)

Amvescap	DSG International	Reed Elsevier	Smith & Nephew
Antofagasta	GlaxoSmithkline	Reuters	Smiths Group
Astrazeneca	Hammerson	Rexam	Standard Chartered
Aviva	Johnson Matthey	Rio Tinto	Tate & Lyle
BAE Systems	Land Securities	Rolls Royce	Tesco
Barclays	Legal & General	Royal Bank of Scotland	United Utilities
BP	Marks & Spencer	Sage Group	Vodafone
British Airways	Morrison Supermarkets	Sainsbury (J)	Whitbread
British Land	NEXT	Scottish & Newcastle	Wolseley
Cadbury Schweppes	Prudential	Scottish Power	WPP Group
Capita Group	Reckitt Benckiser	Slough Estates	

A.8 FTSE 250 Portfolio (Chapter 5)

Aberdeen	Croda International	Marshalls	Spirax Sarco
AGA Foodservice	Daejan Holdings	McAlpine (Alfred)	Spirent
AMEC	De La Rue	Misys	SSL International
Babcock International	Electrocomponents	Morgan Crucible	St James Place Capital
Barret Developments	Euromoney Investor	Morgan Sindall	St Mowdems Property
BBA Group	FKI	National Express	Taylor Woodrow
Bellway	Forth Ports	Northgate	Tomkins
Brixton	GKN	Premier Farnell	Travis Perkins
Brown (N) Group	Grainger Trust	Premier Oil	Trinity Mirror
BSS Group	Greggs	Provident Financial	UK Coal
Bunzl	Hays	Rank Group	United Business Media
Cairn Energy	Headlam Group	Rathbone Brothers	VT Group
Capitol & Regional	Helical Bar	Rentokill Initial	Warner Estates
Carpetright	Hunting	RPS Group	Wetherspoon (JD)
Cattles	IMI	Serco Group	Wilson Bowden
Charter	Interserve	Shanks Group	Wimpey (George)
Chemring	Jardine Lloyd Thompson	SIG	Workspace Group
Close Brothers	Ladbrokes	Signet Group	
Cobham	Laird Group	Smith (DS)	
Crest Nicholson	LogicaCMG	Spectris	

A.9 FTSE Small Cap Portfolio (Chapter 5)

Abacus Group	Castings	Galliford Try	Protherics
Acal	Charles Stanley Group	GCAP Media	Restaurant Group
Alba	Chaucer Holdings	Gleeson (MJ)	Ricardo
Alexon Group	Chime Communications	Hampson Industries	Rok Property Solutions
Amstrad	Chloride Group	Heywood Williams	RPC Group
Anglo Pacific Group	Chrysalis Group	Highway Insurance	Salvesen (Christian)
Anglo-Eastern Plantations	Clarke (T)	Holidaybreak	Simon Group
Anite Group	Clarkson	Hornby	St Ives
Arena Leisure	Creston	Johnson Services	TDG
Ashley (Laura)	Delta	London Scottish Bank	Thorntons
Axis-Shield	Development Securities	Lookers	TT Electronics
Barr (AG)	Devro	Low & Bonar	Uniq
Bespak	Diploma	Management Consulting	VP
Blacks Leisure	Domestic & General	McKay Securities	Wagon
BPP Holdings	DTZ Holdings	Menzies (John)	Whatman
Brammer	Entertainment Rights	Mucklow (A & J)	White Young Green
BPI	Fenner	Nestor Healthcare	WSP Group
Business Post Group	Fisher (James) & Sons	Oxford Instruments	
Care UK	Fuller Smith & Turner	Photo-Me International	

A.10 FTSE Fledgling Portfolio (Chapter 5)

600 Group	City Of London Group	Kewill Systems	Renold
Abbeystead	CML Microsystems	Lambert Howarth	S & U
Air Partner	Cosalt	Lincat Group	Scapa Group
Alphameric	Cropper (James)	Litho Supplies	Tex Holdings
Alumasc	Danka Business Systems	Macfarlane Group	Therataste
Api Group	EDP	Macro 4	Titon Holdings
Armour Group	Ferraris	Mallet	Total Systems
Austin Reed	Fletcher King	Microgen	Trace Group
Avon Rubber	GB Group	Molins	Treatt
Baggeridge Brick	Gibbs & Dandy	Moss Bros	Vega Group
Ben Bailey	Goodwin	MS International	Victoria
Bisichi Mining	Gresham Computing	Northamber	Volex Group
Caffyns	Havelock Europa	Park Group	Waterman Group
Carclo	Haynes Publishing	PGI Group	Windsor
Che Hotel Group	Independent Media	Porvair	

$$\mathbf{A.11} \quad \mathbf{VP(25,1)} = (\mathbf{VP(25)} - \overline{\mathbf{RW}}) \mathbf{e}_1$$

$\mathbf{VP(25)}$ is the variance ratio profile of the $R_1(k)$ statistic $\mathbf{VP(25)} = [R_1(1), R_1(2), \dots, R_1(25)]$.

The first principal component (eigenvector) \mathbf{e}_1 and the mean profile $\overline{\mathbf{RW}}$ are:

Index	\mathbf{e}_1	$\overline{\mathbf{RW}}$	Index	\mathbf{e}_1	$\overline{\mathbf{RW}}$	Index	\mathbf{e}_1	$\overline{\mathbf{RW}}$
1	0.00000	0.00000	11	0.21749	-0.26478	21	0.20378	-0.31917
2	0.13828	-0.14513	12	0.21758	-0.27072	22	0.20109	-0.32509
3	0.16952	-0.18445	13	0.21702	-0.27710	23	0.19856	-0.33046
4	0.18764	-0.19867	14	0.21624	-0.28337	24	0.19635	-0.33546
5	0.19710	-0.20875	15	0.21517	-0.28933	25	0.19390	-0.34009
6	0.20420	-0.21572	16	0.21387	-0.29509			
7	0.20973	-0.22467	17	0.21228	-0.30067			
8	0.21318	-0.23299	18	0.21052	-0.30503			
9	0.21570	-0.24230	19	0.20843	-0.30918			
10	0.21683	-0.25450	20	0.20616	-0.31446			

References

Abhyankar, A., Ghosh, D., Levin, E. and Limmack, R.J. (1997) Bid-ask spreads, trading volume and volatility: intra-day evidence from the London stock exchange, *Journal of Business Finance & Accounting*, 24(3&4), 343–362.

Achelis, S.B. (2001) *Technical analysis from A to Z*, McGraw-Hill, New York, 2nd edition.

Alexander, S.S. (1961) Price movements in speculative markets: trends or random walks, *Industrial Management Review*, 2(2), 7–26.

Alexander, S.S. (1964) Price movements in speculative markets: trends or random walks, number 2, *Industrial Management Review*, 5(2), 25–46.

Amihud, Y. (2002) Illiquidity and stock returns: cross-section and time-series effects, *Journal of Financial Markets*, 5(1), 31–56.

Atanasova, C.V. and Hudson, R.S. (2010) Technical trading rules and calendar anomalies – are they the same phenomena?, *Economics Letters*, 106(2), 128–130.

Barber, B.M., Lee, Y-T., Liu, Y-J. and Odean, T. (2011) The cross-section of speculator skill: evidence from day trading, Working Paper, available at SSRN, ssrn.com/abstract=529063.

Belaire-Franch, J. and Opong, K. (2005) A variance ratio test of the behaviour of some FTSE equity indices using ranks and signs, *Review of Quantitative Finance and Accounting*, 24(1), 93–107.

Bessembinder, H. and Chan, K. (1998) Market efficiency and the returns to technical analysis, *Financial Management*, 27(2), 5–17.

Bessembinder, H. and Venkataraman, K. (2010) Bid-ask spreads: measuring trade execution costs in financial markets, in *Encyclopedia of Quantitative Finance*, R. Cont (ed), John Wiley & Sons.

Biais, B., Hillion, P. and Spatt, C. (1995) An empirical analysis of the limit order book and the order flow in the Paris bourse, *Journal of Finance*, 50(5), 1655–1689.

Bokhari, J., Cai, C., Hudson, R. and Keasey, K. (2005) The predictability and profitability of technical trading rules: does company size matter?, *Economics Letters*, 86(1), 21–27.

Box, G.E.P. and Pierce, D.A. (1970) Distribution of residual autocorrelations in autoregressive-integrated moving average time series models, *Journal of the American Statistical Association*, 65(332), 1509–1526.

Brock, W., Lakonishok, J. and LeBaron, B. (1992) Simple technical trading rules and the stochastic properties of stock returns, *Journal of Finance*, 47(5), 1731–1764.

Brockman, P., Chung, D.Y. and Perignon, C. (2009) Commonality in liquidity: a global perspective, *Journal of Financial and Quantitative Analysis*, 44(4), 851–882.

Brown, S.J., Goetzmann, W.N. and Kumar, A. (1998) The Dow theory: William Peter Hamilton's track record reconsidered, *Journal of Finance*, 53(4), 1311–1333.

Burgess, A.N. (2000) *A computational methodology for modelling the dynamics of statistical arbitrage*, PhD thesis, Department of Decision Sciences, University of London.

Cai, C.X., Hudson, R. and Keasey, K. (2004) Intra day bid-ask spreads, trading volume and volatility: recent empirical evidence from the London stock exchange, *Journal of Business Finance & Accounting*, 31(5&6), 647–676.

Cai, C.X., Hillier, D., Hudson, R. and Keasey, K. (2008) Trading frictions and market structure: an empirical analysis, *Journal of Business Finance & Accounting*, 35(3&4), 563–579.

Cajueiro, D.O. and Tabak, B.M. (2006) Testing for predictability in equity returns for European transition markets, *Economic Systems*, 30(1), 56–78.

Campbell, J.Y., Lo, A.W. and MacKinlay, A.C. (1997) *The econometrics of financial markets*, Princeton University Press, Princeton, New Jersey.

Cecchetti, S.G. and Lam, P-S. (1994) Variance-ratio tests: small-sample properties with an application to international output data, *Journal of Business & Economic Statistics*, 12(2), 177–186.

Chang, E.J., Lima, E.J.A. and Tabak, B.M. (2004) Testing for predictability in emerging equity markets, *Emerging Markets Review*, 5(3), 295–316.

Charles, A. and Darne, O. (2009) Variance-ratio tests of random walk: an overview, *Journal of Economic Surveys*, 23(3), 503–527.

Chelley-Steeley, P. and Park, K. (2008) Evaluating spread decomposition models with a basket security, Paris December 2007 Finance International Meeting AFFI-EUROFIDAI Paper, available at SSRN, ssrn.com/abstract=1070831.

Chen, W.W. and Deo, R.S. (2006) The variance ratio statistic at large horizons, *Econometric Theory*, 22(2), 206–234.

Cheung, Y-W. and Chinn, M.D. (2001) Currency traders and exchange rate dynamics: a survey of the US market, *Journal of International Money and Finance*, 20(4), 439–471.

Cheung, Y-W., Chinn, M.D. and Marsh, I.W. (2004) How do UK-based foreign exchange dealers think their market operates?, *International Journal of Finance and Economics*, 9(4), 289–306.

Choe, K-I., Krausz, J. and Nam, K. (2011) Technical trading rules for nonlinear dynamics of stock returns: evidence from the G-7 stock markets, *Review of Quantitative Finance and Accounting*, 36(3), 323–353.

Chong, T.T-L. and Lam, T-H. (2010) Predictability of nonlinear trading rules in the U.S. stock market, *Quantitative Finance*, 10(9), 1067–1076.

Chordia, T., Roll, R. and Subrahmanyam, A. (2000) Commonality in liquidity, *Journal of Financial Economics*, 56(1), 3–28.

Chordia, T., Sarkar, A. and Subrahmanyam, A. (2005) An empirical analysis of stock and bond market liquidity, *Review of Financial Studies*, 18(1), 85–129.

Chow, K.V. and Denning, K.C. (1993) A simple multiple variance ratio test, *Journal of Econometrics*, 58(3), 385–401.

Christie, W.G. and Schultz, P.H. (1994) Why do NASDAQ market makers avoid odd-eighth quotes?, *Journal of Finance*, 49(5), 1813–1840.

Christie, W.G., Harris, J.H. and Schultz, P.H. (1994) Why did NASDAQ market makers stop avoiding odd-eighth quotes?, *Journal of Finance*, 49(5), 1841–1860.

Dacorogna, M.M., Gencay, R., Muller, U., Olsen, R.B. and Pictet, O.V. (2001) *An introduction to high-frequency finance*, Academic Press, San Diego, California.

Dawson, E.R. and Steeley, J.M. (2003) On the existence of visual technical patterns in the UK stock market, *Journal of Business Finance & Accounting*, 30(1/2), 263–293.

Day, T.E. and Wang, P. (2002) Dividends, nonsynchronous prices, and the returns from trading the Dow Jones Industrial Average, *Journal of Empirical Finance*, 9(4), 431–454.

De Jong, F. and Rindi, B. (2009) *The microstructure of financial markets*, Cambridge University Press, Cambridge.

Dempster, M.A.H. and Jones, C.M. (2002) Can channel pattern trading be profitably automated?, *European Journal of Finance*, 8(3), 275–301.

Easley, D. and O'Hara, M. (1987) Price, trade size, and information in securities markets, *Journal of Financial Economics*, 19(1), 69–90.

Easley, D., Kiefer, N.M. and O'Hara, M. (1997) The information content of the trading process, *Journal of Empirical Finance*, 4(2/3), 159–186.

Efron, B. (1979) Bootstrap methods: another look at the jackknife, *Annals of Statistics*, 7(1), 1–26.

Efron, B. and Tibshirani, R.J. (1998) *An introduction to the bootstrap*, Chapman & Hall/CRC Monographs on Statistics and Applied Probability, Boca Raton, Florida.

Ellis, C.A. and Parbery, S.A. (2005) Is smarter better? A comparison of adaptive, and simple moving average trading strategies, *Research in International Business and Finance*, 19(3), 399–411.

Enders, W. and Granger, C.W.J. (1998) Unit-root tests and asymmetric adjustment with an example using the term structure of interest rates, *Journal of Business & Economic Statistics*, 16(3), 304–311.

Fama, E.F. and Blume, M.E. (1966) Filter rules and stock market trading, *Journal of Business*, 39(1:2), 226–241.

Fama, E.F. (1970) Efficient capital markets: a review of theory and empirical work, *Journal of Finance*, 25(2), 383–417.

Fama, E.F. (1991) Efficient capital markets: II, *Journal of Finance*, 46(5), 1575–1617.

Fang, Y. and Xu, D. (2003) The predictability of asset returns: an approach combining technical analysis and time series forecasts, *International Journal of Forecasting*, 19(3), 369–385.

Gehrig, T. and Menkhoff, L. (2004) The use of flow analysis in foreign exchange: exploratory evidence, *Journal of International Money and Finance*, 23(4), 573–594.

Gehrig, T. and Menkhoff, L. (2006) Extended evidence on the use of technical analysis in foreign exchange, *International Journal of Finance and Economics*, 11(4), 327–338.

Gencay, R. (1998) The predictability of security returns with simple technical trading rules, *Journal of Empirical Finance*, 5(4), 347–359.

Gencay, R. and Stengos, T. (1998) Moving average rules, volume and the predictability of security returns with feedforward networks, *Journal of Forecasting*, 17(5/6), 401–414.

George, T.J., Kaul, G. and Nimalendran, M. (1991) Estimation of the bid-ask spread and its components: a new approach, *Review of Financial Studies*, 4(4), 623–656.

Glosten, L.R. and Milgrom, P.R. (1985) Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics*, 14(1), 71–100.

Glosten, L.R. and Harris, L.E. (1988) Estimating the components of the bid/ask spread, *Journal of Financial Economics*, 21(1), 123–142.

Goldberg, D.E. (1989) *Genetic algorithms in search, machine learning and optimization*, Addison-Wesley, Reading, Massachusetts.

Goyenko, R.Y. and Ukhov, A.D. (2009) Stock and bond market liquidity: a long-run empirical analysis, *Journal of Financial and Quantitative Analysis*, 44(1), 189–212.

Goyenko, R.Y., Holden, C.W. and Trzcinka, C.A. (2009) Do liquidity measures measure liquidity?, *Journal of Financial Economics*, 92(2), 153–181.

Grossman, S.J. and Stiglitz, J.E. (1980) On the impossibility of informationally efficient markets, *American Economic Review*, 70(3), 393–408.

Hasbrouck, J. (1991) Measuring the information content of stock trades, *Journal of Finance*, 46(1), 179–207.

Hasbrouck, J. (2007) *Empirical market microstructure: the institutions, economics, and econometrics of securities trading*, Oxford University Press, New York.

Hoque, H.A.A.B., Kim, J.H. and Pyun, C.S. (2007) A comparison of variance ratio tests of random walk: a case of Asian emerging markets, *International Review of Economics and Finance*, 16(4), 488–502.

Huang, R.D and Stoll, H.R. (1997) The components of the bid-ask spread: a general approach, *Review of Financial Studies*, 10(4), 995–103.

Hudson, R., Dempsey, M. and Keasey, K. (1996) A note on the weak form efficiency of capital markets: the application of simple technical trading rules to UK stock prices – 1935 to 1994, *Journal of Banking & Finance*, 20(6), 1121–1132.

Hung, J-C., Lee, Y-H. and Pai, T-Y. (2009) Examining market efficiency for large- and small-capitalization of TOPIX and FTSE stock indices, *Applied Financial Economics*, 19(9), 735–744.

Ito, M. and Sugiyama, S. (2009) Measuring the degree of time varying market inefficiency, *Economics Letters*, 103(1), 62–64.

Jensen, M.C. and Benington, G.A. (1970) Random walks and technical theories: some additional evidence, *Journal of Finance*, 25(2), 469–482.

Jensen, M.C. (1978), Some anomalous evidence regarding market efficiency, *Journal of Financial Economics*, 6(2&3), 95–101.

Jolliffe, I.T. (2002) *Principal component analysis*, Springer, New York, 2nd edition.

Kaufman, P.J. (2005) *New trading systems and methods*, John Wiley & Sons, New York, 4th edition.

Kim, J.H. (2006) Wild bootstrapping variance ratio tests, *Economics Letters*, 92(1), 38–43.

Kim, J.H. and Shamsuddin, A. (2008) Are Asian stock markets efficient? Evidence from new multiple variance ratio tests, *Journal of Empirical Finance*, 15(3), 518–532.

Korajczyk, R.A. and Sadka, R. (2008) Pricing the commonality across alternative measures of liquidity, *Journal of Financial Economics*, 87(1), 45–72.

Kwon, K-Y. and Kish, R.J. (2002) A comparative study of technical trading strategies and return predictability: an extension of Brock, Lakonishok and LeBaron (1992) using NYSE and NASDAQ indices, *Quarterly Review of Economics and Finance*, 42(3), 611–631.

Kyle, A.S. (1985) Continuous auctions and insider trading, *Econometrica*, 53(6), 1315–1335.

Leigh, W., Frolich, C.J., Hornik, S., Purvis, R.L. and Roberts, T.L. (2008) Trading with a stock chart heuristic, *IEEE Transactions on Systems, Man, and Cybernetics Part A*, 38(1), 93–104.

Lesmond, D.A. (2005) Liquidity of emerging markets, *Journal of Financial Economics*, 77(1), 411–452.

Lim, K-P. (2007) Ranking market efficiency for stock markets: a nonlinear perspective, *Physica A*, 376, 445–454.

Lim, K-P. and Brooks, R. (2011) The evolution of stock market efficiency over time: a survey of the empirical literature, *Journal of Economic Surveys*, 25(1), 69–108.

Lin, J-C., Sanger, G.C. and Booth, G.G. (1995) Trade size and components of the bid-ask spread, *Review of Financial Studies*, 8(4), 1153–1183.

Lindemann, A., Dunis, C.L. and Lisboa, P. (2005) Extending the variance ratio test to visualize structure in data: an application to the S&P 100 Index, *Applied Financial Economics Letters*, 1(3), 189–197.

Ljung, G.M. and Box, G.E.P. (1978) On a measure of lack of fit in time series models, *Biometrika*, 65(2), 297–303.

Lo, A.W. and MacKinlay, A.C. (1988) Stock market prices do not follow random walks: evidence from a simple specification test, *Review of Financial Studies*, 1(1), 41–66.

Lo, A.W. and MacKinlay, A.C. (1989) The size and power of the variance ratio test in finite samples: a Monte Carlo investigation, *Journal of Econometrics*, 40(2), 203–238.

Lo, A.W., Mamaysky, H. and Wang, J. (2000) Foundations of technical analysis: computational algorithms, statistical inference and empirical implementation, *Journal of Finance*, 55(4), 1705–1765.

Lo, A.W. (2004) The adaptive market hypothesis: market efficiency from an evolutionary perspective, *Journal of Portfolio Management*, 30(5), 15–29.

Lui, Y-H. and Mole, D. (1998) The use of fundamental and technical analyses by foreign exchange dealers: Hong Kong evidence, *Journal of International Money and Finance*, 17(3), 535–545.

Madhavan, A., Richardson, M. and Roomans, M. (1997) Why do security prices change? A transaction-level analysis of NYSE stocks, *Review of Financial Studies*, 10(4), 1035–1064.

Makridakis, S.G., Wheelwright, S.C. and Hyndman, R.J. (1998) *Forecasting: methods and applications*, John Wiley & Sons, New York, 3rd edition.

Malkiel, B.G. (1996) *A random walk down Wall Street*, W.W. Norton, New York, 6th edition.

Marshall, B.R. and Cahan, R.H. (2005) Is technical analysis profitable on a stock market which has characteristics that suggest it may be inefficient?, *Research in International Business and Finance*, 19(3), 384–398.

Marshall, B.R., Qian, S. and Young, M. (2009) Is technical analysis profitable on US stocks with certain size, liquidity or industry characteristics?, *Applied Financial Economics*, 19(15), 1213–1221.

Masters, T. (1998) Just what are we optimizing, anyway?, *International Journal of Forecasting*, 14(2), 277–290.

Menkhoff, L. (1997) Examining the use of technical currency analysis, *International Journal of Financial Economics*, 2(4), 307–318.

Menkhoff, L. (2010) The use of technical analysis by fund managers: international evidence, *Journal of Banking & Finance*, 34(11), 2573–2586.

Mills, T.C. (1997) Technical analysis and the London stock exchange: testing trading rules using the FT30, *International Journal of Financial Economics*, 2(4), 319–331.

Mitchell, M. (1998) *An introduction to genetic algorithms*, MIT Press, Cambridge, Massachusetts.

Nam, K., Washer, K.M. and Chu, Q.C. (2005) Asymmetric return dynamics and technical trading strategies, *Journal of Banking & Finance*, 29(2), 391–418.

Neal, R. and Wheatley, S.M. (1998) Adverse selection and bid-ask spreads: evidence from closed-end funds, *Journal of Financial Markets*, 1(1), 121–149.

O'Hara, M. (2003) Presidential address: liquidity and price discovery, *Journal of Finance*, 58(4), 1335–1354.

Oberlechner, T. (2001) Importance of technical and fundamental analysis in the European foreign exchange market, *International Journal of Financial Economics*, 6(1), 81–93.

Olson, D. (2004) Have trading rule profits in the currency markets declined over time?, *Journal of Banking & Finance*, 28(1), 85–105.

Park, C-H. and Irwin, S.H. (2007) What do we know about the profitability of technical analysis?, *Journal of Economic Surveys*, 21(4), 786–826.

Phengpis, C. (2006) Are emerging stock market price indices really stationary?, *Applied Financial Economics*, 16(13), 931–939.

Press, W.H., Teukolsky, S.A., Vetterling, W.T. and Flannery, B.P. (1996) *Numerical recipes in C: the art of scientific computing*, Cambridge University Press, New York, 2nd edition.

Pring, M.J. (2002) *Technical analysis explained: the successful investor's guide to spotting investment trends and turning points*, McGraw-Hill, New York, 4th edition.

Puckett, A. and Yan, X.S. (2011) The interim trading skills of institutional investors, *Journal of Finance*, 66(2), 601–633.

Roll, R. (1984) A simple implicit measure of the effective bid-ask spread in an efficient market, *Journal of Finance*, 39(4), 1127–1139.

Self, J.K. and Mathur, I. (2006) Asymmetric stationarity in national stock market indices: an MTAR analysis, *Journal of Business*, 79(6), 3153–3174.

Sullivan, R., Timmermann, A. and White, H. (1999) Data-snooping, technical trading rule performance, and the bootstrap, *Journal of Finance*, 54(5), 1647–1691.

Tabak, B.M. (2003) The random walk hypothesis and the behaviour of foreign capital portfolio flows: the Brazilian stock market case, *Applied Financial Economics*, 13(5), 369–378.

Takagi, T. and Sugeno, M. (1985) Fuzzy identification of systems and its applications to modeling and control, *IEEE Transactions On Systems Man And Cybernetics*, 15(1), 116–132.

Taylor, M.P. and Allen, H. (1992) The use of technical analysis in the foreign exchange market, *Journal of International Money and Finance*, 11(3), 304–314.

Taylor, S.J. (2005) *Asset price dynamics, volatility and prediction*, Princeton University Press, Princeton, New Jersey.

Tsay, R.S. (2010) *Analysis of financial time series*, John Wiley & Sons, Hoboken, New Jersey, 3rd edition.

Van Horne, J.C. and Parker, G.G.C. (1967) The random-walk theory: an empirical test, *Financial Analysts Journal*, 23(6), 87–92.

Van Horne, J.C. and Parker, G.G.C. (1968) Technical trading rules: a comment, *Financial Analysts Journal*, 24(4), 128–132.

Van Ness, B.F., Van Ness, R.A. and Warr, R.S. (2001) How well do adverse selection components measure adverse selection?, *Financial Management*, 30(3), 77–98.

Wang, J-L. and Chan, S-H. (2007) Stock market trading rule discovery using pattern recognition and technical analysis, *Expert Systems with Applications*, 33(2), 304–315.

Wang, J-L. and Chan, S-H. (2009) Trading rule discovery in the US stock market: an empirical study, *Expert Systems with Applications*, 36(3), 5450–5455.

Welch, I. and Goyal, A. (2008) A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies*, 21(4), 1455–1508.

Whang, Y-J. and Kim, J. (2003) A multiple variance ratio test using subsampling, *Economics Letters*, 79(2), 225–230.

White, H. (1980) A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica*, 48(4), 817–838.

Wong, W-K., Manzur, M. and Chew, B-K. (2003) How rewarding is technical analysis? Evidence from Singapore stock market, *Applied Financial Economics*, 13(7), 543–551.

Wright, J.H. (2000) Alternative variance-ratio tests using ranks and signs, *Journal of Business & Economic Statistics*, 18(1), 1–9.