

NOTE TO USERS

The original document received by UMI contains pages with slanted print. Pages were microfilmed as received.

This reproduction is the best copy available

UMI

THE UNIVERSITY OF CALGARY

Predictability of Energy Futures Prices

by

Sean McGrath

A THESIS

**SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE
DEGREE OF MASTER OF ARTS**

DEPARTMENT OF ECONOMICS

CALGARY, ALBERTA

APRIL, 1998

©Sean McGrath 1998



**National Library
of Canada**

**Acquisitions and
Bibliographic Services**

**395 Wellington Street
Ottawa ON K1A 0N4
Canada**

**Bibliothèque nationale
du Canada**

**Acquisitions et
services bibliographiques**

**395, rue Wellington
Ottawa ON K1A 0N4
Canada**

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-31299-2

Canada

ABSTRACT

This paper examines the debate over the predictability of asset returns. Specifically, the efficient markets approach is compared with that of technical analysis, within the context of returns in energy futures markets. It is found that, due to very different starting assumptions, a comparison of the two is difficult. While statistical deviations from the random walk exist, it is not clear that these represent an economic deviation from efficient markets.

ACKNOWLEDGEMENTS

I would like to thank Dr. Serletis for his guidance and encouragement in helping me see this thesis through to completion. I would also like to thank the many friends and family members who provided encouragement and support throughout this whole endeavor. Finally, a special thanks to Eimer and Kevin, who listened to many an earful as I muddled my way through this program.

DEDICATION

To my Mom.

TABLE OF CONTENTS

Approval Page.....	ii
Abstract.....	iii
Acknowledgements.....	iv
Dedication.....	v
Table of Contents.....	vi
List of Charts.....	x
List of Tables.....	xi
CHAPTER ONE: INTRODUCTION.....	1
1.1 Motivation for the Study.....	1
1.2 Outline of the Study.....	2
1.3 Objectives of the Study.....	4
CHAPTER TWO: BACKGROUND THEORY.....	5
2.1 Introduction.....	5
2.2 Technical Analysis.....	6
2.2.1 Origins.....	6
2.2.2 Philosophy.....	9
2.2.3 Theory.....	14
2.2.4 Summary.....	25
2.3 Efficient Markets Theory.....	26
2.3.1 Origins.....	26
2.3.2 Philosophy.....	27

2.3.3 Random Walk Hypotheses.....	30
2.3.4 Implications.....	33
2.4 Behavioural Finance.....	35
2.4.1 Origins and Philosophy of Behavioural Finance.....	35
2.4.2 Implications.....	37
2.5 Conclusion.....	38
CHAPTER THREE: RECENT EMPIRICAL EVIDENCE.....	40
3.1 Introduction.....	40
3.2 Efficient Markets Theory.....	41
3.2.1 Tests for Return Predictability.....	43
3.2.2 Event Studies.....	48
3.2.3 Tests for Private Information.....	50
3.3 Technical Analysis.....	52
3.3.1 Neftci (1991).....	53
3.3.2 Taylor and Allen (1992).....	55
3.3.3 Brock, Lakonishok and LeBaron (1992).....	56
3.3.4 Blume, Easley and O'Hara (1994).....	57
3.3.5 Taylor (1994).....	58
3.3.6 Studies in <i>Technical Analysis of Stocks and Commodities</i> Magazine.....	59
3.4 Energy Futures Markets.....	60
3.4.1 Gjolberg (1985).....	61
3.4.2 Bopp and Sitzer (1988).....	62

3.4.3 Panas (1990).....	63
3.4.4 Serletis (1992).....	63
3.4.5 Nainar (1993).....	64
3.4.6 Herbert (1995).....	65
3.5 Conclusion.....	66
CHAPTER FOUR: STATISTICAL METHODOLOGY.....	69
4.1 Introduction.....	69
4.2 Tests of Random Walk 1.....	70
4.2.1 Autocorrelation Coefficients.....	71
4.2.2 Ljung-Box Q-Statistic.....	71
4.3 Tests of Random Walk 2.....	72
4.3.1 Moving Average Crossover.....	73
4.3.2 Channel Rule.....	75
4.4 Tests of Random Walk 3.....	76
4.4.1 Variance Ratios.....	77
4.5 Unit Root Tests.....	77
4.5.1 Augmented Dickey-Fuller (ADF) Test.....	78
4.6 Summary.....	80
CHAPTER FIVE: DATA AND RESULTS.....	82
5.1 Introduction.....	82
5.2 Data.....	82
5.3 Tests of RW1.....	87
5.4 Tests of RW2 – Results of Technical Trading Rules.....	89

5.4.1 Commodity Channel Index (CCI).....	89
5.4.2 Moving Average (MA) Crossover with Optimization.....	97
5.5 Tests of RW3.....	99
5.6 Unit Root Tests.....	101
5.7 Summary.....	104
CHAPTER SIX: CONCLUSIONS.....	106
6.1 Commenting on the Paradox.....	106
6.2 Reconciling the Theories.....	108
6.3 Directions for Future Research.....	110
REFERENCES.....	112

LIST OF CHARTS

Figure 2.1: Trendlines and Channels.....	16
Figure 2.2: Head and Shoulders Reversal Pattern.....	22
Chart 5.1 – Logged Daily Closing Prices for Crude Oil Spot Contracts.....	84
Chart 5.2 – Logged Daily Closing Prices for Heating Oil Spot Contracts.....	84
Chart 5.3 – Logged Daily Closing Prices for Liquid Propane Spot Contracts.....	85
Chart 5.4 – Logged Daily Closing Prices for Natural Gas Spot Contracts.....	85
Chart 5.5 – Logged Daily Closing Prices for Electricity Spot Contracts.....	86
Chart 5.6 – Logged Daily Closing Prices for Unleaded Gasoline Spot Contracts ...	86

LIST OF TABLES

Table 2.1: Classification of Random Walk and Martingale Hypotheses..... 31

Table 5.1: Details Regarding the Data..... 83

Table 5.2 – Summary statistics for daily percentage returns in energy futures..... 87

Table 5.3 – Autocorrelation in daily percentage returns in energy futures..... 88

Table 5.4 – CCI Trading Rule Results for Liquid Propane..... 91

Table 5.5 – CCI Trading Rule Results for Natural Gas..... 92

Table 5.6 – CCI Trading Rule Results for Crude Oil..... 93

Table 5.7 – CCI Trading Rule Results for Heating Oil..... 94

Table 5.8 – CCI Trading Rule Results for Palo Verde Electricity Futures..... 95

Table 5.9 – CCI Trading Rule Results for Unleaded Gasoline Futures..... 96

Table 5.10 – MA Crossover Results for Unleaded Gasoline..... 97

Table 5.11 – MA Crossover Results for Palo Verde Electricity..... 97

Table 5.12 – MA Crossover Results for Natural Gas..... 97

Table 5.13 – MA Crossover Results for Liquid Propane..... 98

Table 5.14 – MA Crossover Results for Heating Oil..... 98

Table 5.15 – MA Crossover Results for Crude Oil..... 98

Table 5.16 – Variance Ratios for Percentage Daily Returns in Energy Futures..... 100

Table 5.17 – Tests for a unit root in the logarithms of the series..... 102

Table 5.18 – Tests for a unit root in the first differences of the logarithms
of the series..... 103

CHAPTER 1 – INTRODUCTION

This thesis addresses one of the great paradoxes of modern economics. Specifically, this study explores the debate between financial economists and the investment industry over the predictability of asset returns. The two schools of thought reach starkly different conclusions on the subject, yet practitioners in each school seem to reject the evidence provided by the other. This thesis aims at commenting on this paradox, within the specific context of the predictability of energy futures returns using technical analysis.

This introduction will continue as follows. In the first section, the motivation for this study is examined more thoroughly. The second section provides a brief outline of the study itself. Finally, the third section outlines the objectives of the study.

1.1 Motivation for the Study

The primary motivation for this study is to comment on, or partially resolve, the paradox mentioned above. The fact that financial economists provide seemingly irrefutable evidence in favor of efficient markets would seem to be irreconcilable with the fact that the majority of people are willing to pay for financial advice. Neither school of thought seems to capture the *whole* truth. Financial economists view economic agents as rational, and see markets as efficient. Yet, those same rational agents are observed as being willing to pay for advice that would enable them to exploit an *inefficient* market.

Interesting results are presented in Taylor & Allen (1992). Within the context of the London foreign exchange markets, the survey shows that the vast majority of traders

rely largely on predictions based on technical analysis. Specifically, 90% of traders rely on technical analysis over shorter time horizons, from intraday to one-week predictions. Furthermore, over longer time horizons, traders rely heavily on fundamental analysis. The key here is that traders believe that these markets *are* predictable based on publicly available information. Essentially, they believe that these markets are inefficient. In addition, Taylor & Allen (1990) provide evidence that the predictions of technical analysts outperform predictions based on the random walk, vector autoregression and univariate autoregressive moving-average time-series models.

Therefore, we have evidence not only that investors are willing to pay for analysts' predictions, but also that traders, acting in their own self interest, act with the belief that markets are inefficient. Clearly, any plausible version of the rational expectations or efficient markets hypotheses will need to account for this sort of behavior. This study will proceed to examine the two approaches being taken – those of the financial economist and of the technical analyst – and attempt to comment on this paradox based on the evidence.

1.2 Outline of the Study

This introductory chapter is followed by an examination of the background theory in Chapter 2. Given the fact that the two schools of thought have evolved along very different lines, a thorough examination of the background theory is all the more important. Section 2.2 outlines the basic ideas of technical analysis, including its origins, underlying philosophy and theory. Section 2.3 turns to the efficient markets theory, again examining its origins, underlying philosophy and theory. The emerging field of

behavioral finance is introduced in Section 2.4 as a possible reconciliation of the previous two theories.

Chapter 3 provides an overview of some of the recent empirical evidence. In fact, there have been surprisingly few formal studies of technical analysis, and still fewer of technical analysis in the energy futures markets. Therefore, the chapter adopts a rather piecemeal approach, summarizing the results that do relate to these topics. Section 3.2 reviews the results concerning the efficient markets theory, including tests of return predictability, event studies and tests for private information. Section 3.3 looks at the results of studies on technical analysis, while Section 3.4 looks at the evidence in energy futures markets.

The statistical methodology to be used in the study is outlined in Chapter 4. Tests of the various versions of the random walk hypothesis are outlined in sections 4.2 through 4.4. Under the different assumptions of the random walk, appropriate tests include autocorrelation coefficients, Q statistics and variance ratios as well as tests of technical trading rules, such as the moving average crossover and the channel rule. Unit root tests are outlined in Section 4.5.

Chapter 5 presents the data and results of the study. Section 5.2 summarizes and plots the data, including specific details of the energy futures contracts to be used. Sections 5.3 through 5.6 report the results of those tests outlined in the previous chapter on methodology. Finally, Section 5.7 provides an interpretation of the results.

The conclusions of the thesis are presented in Chapter 6. The entirety of the study is drawn upon in an attempt to comment on the debate between financial economists and the investment industry. Directions and topics for future research are suggested.

1.3 Objectives of the Study

The primary objective of this study is to comment on the paradox mentioned above. Clearly, the two schools of thought – those of financial economists and technical analysts – are to some degree incompatible. The issues of market efficiency and rational expectations are of particular interest. Some modification of either or both theories is needed, in order to account for the evidence of the other.

The author is under no illusion that a paradox of this magnitude is to be resolved. It is for this reason that the objective of this study is only to *comment* on the paradox. It is felt that a close examination of the underlying theory and assumptions of both schools of thought should serve to diminish the paradox somewhat. That is, when certain assumptions are relaxed, it can hopefully be shown that the two schools in fact share much common ground.

In reconciling the competing theories somewhat, it may also be possible to point out some fruitful areas for future research. That is, to the extent that the two schools of thought can be shown to share common ground, ideas of one school may point the way towards new directions in the other. The point of view taken in this study is that of the financial economist, and therefore an additional objective is to point out ways in which the ideas of technical analysis may be incorporated into future academic work.

CHAPTER 2 - BACKGROUND THEORY

2.1 Introduction

This chapter provides an overview of the background theory needed in order to proceed with a formal study of technical analysis. Essentially, a study of technical analysis implies a study of the efficient markets literature. The very notion that technical analysis - a study of past market behavior and therefore of publicly available information - can be used to outperform the market would seem to be inconsistent with an efficient market. In addition, technical analysis implies that market participants make decisions in a way which is at times irrational, which would again be inconsistent with the efficient markets literature. A thorough understanding of this background theory becomes all the more important when it is realized that the academic and non-academic literature on the subject seem to have evolved in two very different ways, and at times seems to be written in two different languages.

The chapter will be organized as follows. Section 2.2 will outline the theories and trading rules known collectively as *technical analysis*. Section 2.3 will outline the key concepts associated with the *efficient markets theory*. Recent theories of investor behavior, from the realm of *behavioral finance*, will be examined in Section 2.4. Finally, in Section 2.5 an attempt is made to reconcile some of these theories, or at least to articulate the competing theories in a common language.

2.2 Technical Analysis

Murphy (1986, p.1), an authoritative source in present-day technical analysis, defines the subject:

Technical analysis is the study of market action, primarily through the use of charts, for the purpose of forecasting future price trends. The term “market action” includes the three principal sources of information available to the technician - price, volume and open interest.

Murphy goes on to emphasize that proper technical analysis involves all of the above variables. The price trends and patterns that the technical analyst sees must be accompanied by corresponding trends in volume and / or open interest in order to be valid. As will be shown in later sections, tests of weak-form efficiency often consider past price data but ignore volume, although this is clearly not consistent with the basic premise of technical analysis.

2.2.1 Origins

Technical analysis dates back to the 1800s, if not earlier, and is often seen as the original form of investment analysis. Brock, Lakonishok and LeBaron (1992) point out that in these early days of the stock market, the lack of disclosure of financial information precluded the development of fundamental analysis. In the absence of reliable information, and given relatively primitive computing capabilities, technical analysis developed as a practical means of attempting to forecast market movements.

The majority of the ideas within technical analysis find their origin in the Dow Theory - thus Edwards and Magee (1966, p.11), perhaps *the* authoritative source on the subject, state that “the Dow Theory is the granddaddy of all technical market studies”. The Dow Theory, as it came to be known, was in fact a posthumous collection of the ideas and writings of Charles H. Dow, founder of the Dow-Jones financial news service. In a series of articles published in the *Wall Street Journal* in the late 1800s, Dow outlined the framework of his theory on stock price movements.

The starting point for Dow involved the creation in 1897 of what were the two original market averages - the Dow-Jones Industrials and the Dow-Jones Rails. In 1929 a separate Utilities average was created, and of course since that time innumerable market averages and indexes have been created, each aimed at tracking a certain sector of the overall market. The basic premise behind creating these averages was to identify overall market trends. It was considered common knowledge that securities prices tended to move together, and that these movements consisted of primary uptrends known as *bull markets*, and primary downtrends known as *bear markets*. The Dow Theory essentially aimed at identifying market trend as quickly and reliably as possible.

Dow further refined this notion of trend, specifying three types of trend. The *primary trend*, as outlined above, described the longer-term overall market trend, lasting usually a year or longer. The *secondary trend* was seen as a reaction or correction of the primary trend, moving in the opposite direction for a period of weeks or months. Finally, *minor trends* were seen as daily fluctuations in prices which were generally unpredictable but unimportant in terms of the larger trends at work. The key to investing, in this framework, lay simply in properly identifying the trend.

Investor behavior was seen as a combination of far-sighted investors as well as short-sighted, naive investors. Within a primary trend, these two groups would act in different ways. For example, in a bull market the first phase is accumulation. This is a time when markets are at their lows - all the bad news is out, the “public” has lost faith in the market, but sophisticated investors see this and begin to accumulate shares. Eventually, due to this accumulation, prices begin to rise slowly. Finally, after a long period of price increases, “public” buying comes into play and prices rise at a dramatic rate. The naive investor, having missed out on the majority of the price increases, now decides to buy. This phase actually corresponds to the distribution phase, in which the far-sighted investor begins to sell shares and realize significant gains. This selling pressure eventually overcomes the euphoric buying, and price declines bring on a phase of panic selling. When the selling eventually subsides, the market consolidates (moving sideways) and eventually moves into the accumulation phase again.

Dow’s theory involved predicting these primary trends by means of the *principle of confirmation*. Simply put, the averages must confirm one another in terms of trend. If the Industrials reached new highs on a rally, while a rally in the Rails failed to reach new highs, the averages were said to show a *bearish divergence*. A divergence of the sort was seen as a sign of a changing trend. Another aspect of *confirmation* was that volume was required to confirm a trend direction - a price movement without corresponding volume increase was seen as suspect.

Dow’s contribution to the overall theory of technical analysis cannot be overestimated. While technical trading rules have evolved considerably, and while Dow’s rules may seem simplistic by modern standards, the pervasiveness of his ideas has

only increased over the years. Business headlines are preoccupied with the levels of the Dow-Jones Industrials and other averages, and modern trading rules still base themselves on the notions of trend and confirmation. Dow's original ideas have vaulted the theory of technical analysis beyond the realm of its practitioners, to the point where it is entrenched in the public imagination and folklore.

2.2.2 Philosophy

Murphy (1986 p.2) outlines the three basic premises upon which technical analysis is founded:

1. Market action discounts everything.
2. Prices move in trends.
3. History repeats itself.

Again, this is essentially a reformulation of Dow's ideas, but Murphy presents a more concise argument. First, consider the statement "market action discounts everything". Basically, this implies that anything which can affect the price of a security is reflected in its price already - thus the focus on past market action alone. If price and volume has been rising, the technician sees this as evidence of an uptrend. The reasons for this market action are themselves unimportant. Whether it is a political event which may affect the price of oil, or a rumored management change in a given company, the net effect is that any information that will affect price will do just that. In fact, the technician is arguing that markets are efficient in terms of information regarding the valuation of prices, and prices will only fluctuate as a result of changing levels of supply and demand. Clearly, within this point of view is an inherent belief that people interpret information

differently, or that they buy and sell securities based on reasons other than information regarding valuation.

Next, consider the statement “prices move in trends”. Directly from Dow Theory, the notion that prices trend is also essential to technical analysis. The above interpretation of prices resulting from supply and demand would reduce to a trivial case if it were not for the concept of trend. It is the longer-term imbalances between supply and demand which in fact *cause* trends to occur. Again using the logic of the two investor types - the sophisticated and the naive investor - we can interpret the nature of these imbalances. For example, following a market decline we enter a period of market consolidation, in which prices neither increase or decrease substantially but rather trade within a given price range. During this phase the naive investor, who has likely suffered substantial losses after having bought near the top, becomes disgruntled and finally sells. The sophisticated investor, having sold his positions near the market top, now begins patiently to accumulate shares at the depressed prices. There is a relative *balance* of supply and demand as shares turn over between sophisticated and naive investors. When the sellers begin to thin out and with prices still at low levels, accumulation remains attractive for the sophisticated investor, and an imbalance arises causing excess demand for shares. Due to this excess demand, an upward trend ensues, punctuated by the secondary trends in which some short-horizon sophisticated investors begin to take profits. Finally, the same dynamics are said to take place at a market top, with the sophisticated investor in a distribution phase. Clearly there is, inherent in this line of reasoning, the belief that investors are heterogeneous in nature, with the naive investor (meaning small, uninformed) losing out to the sophisticated investor. Technical analysis,

it is argued, allows one to observe these trends as they emerge, and to profitably join the trend with the sophisticated investors.

Finally, consider the statement “history repeats itself”. This simply means that the technical analyst believes that the same dynamics which have moved markets in the past will continue to be the dynamics which move markets in the future. In its most basic form, this implies that markets will continue to trend and that every bull market will inevitably be followed by a bear market, and vice-versa. More specifically, in terms of the dynamics of the supply and demand imbalance, this implies that *patterns* observed in charts in the past may shed light on future price movements, should those same patterns occur again. The types of patterns that the technical analyst looks for will be examined in the following section. The important point here is that these patterns result from the interplay between supply and demand - the physical turnover of shares between buyers and sellers - and the dynamics of this interplay are assumed to recur over time.

The philosophical basis of technical analysis, as outlined above, opens the door to several related questions, as well as plentiful criticism. Perhaps the most common criticism leveled at technical analysis concerns the notion that “history repeats itself”. The efficient markets theorists will argue that, given “markets discount everything” in an efficient manner, expected price moves going forward *must* be random in nature. Therefore, the past cannot be used to predict the future. However, it should be noted that the statement “history repeats itself” is nothing other than the basis of inductive reasoning. Thus, Murphy (1986, p.19) asks, in responding to the criticism that past data cannot be used to predict future price moves, “what other data is there?” In other words, even the efficient markets / random walk hypothesis depends on history repeating itself,

in the sense that we must assume that the market will *continue* to act in an efficient manner. This issue will be explored in greater detail in a later section, in which the philosophical and logical foundations of the efficient markets / random walk theory are examined.

The simplicity of technical analysis, again as outlined in the premises above, suggests several extensions. Given that the technical analyst believes that prices trend due to supply and demand imbalances (etc.), his analysis should be applicable to a broad array (perhaps any) price series which results from trading in an open market. Indeed, this is one of the claims of technical analysts - they argue that their methods are applicable to stocks, bonds, commodities, futures and currencies. Another extension involves the claim that the same dynamics which move prices over longer horizons should be seen over shorter horizons, even intraday. Simply put, if trends and patterns result from supply and demand interaction, this interaction should in fact define the microstructure of the market. While this is not in line with earlier theory that minor trends are unpredictable and unimportant, it is a logical extension given modern computing abilities, improved data availability and the widespread existence of intraday trading.

Another issue involves the interplay between technical and fundamental analysis. Fundamental analysis involves the study of the *valuation* of a company, using a vast array of ratios and indicators such as the price-earnings ratio, price-cashflow, price-book value, etc. Undervalued stocks should be bought, while overvalued stocks should be sold. The technician, however, would see this exercise as utterly futile, based on the notion that “markets discount everything”. Therefore it would be misguided, as well as arrogant, to

believe that one's own valuation of a company was in some sense more accurate than the overall market's valuation. More importantly, past price and volume movements, should *reveal* whether or not a consensus is emerging that a given company is undervalued or overvalued. After all, the technician would argue, it is irrelevant whether or not a certain stock *is* undervalued - the only issue is that there be a growing consensus regarding the supposed undervaluation, which generates excess demand, and therefore moves the price. Thus, the pure technician sees no role for fundamental analysis.

Underlying the entire foundation of technical analysis is the idea that it is an *art*. By this it is meant that trading rules are never to be followed blindly or mechanically. The analyst is always looking for clues which should confirm or make him suspicious of his interpretation of his charts. The pure technician would look for these *clues* solely within the charts (size and length of previous trend, false moves, etc.), while most modern technicians will seek clues in related technical indicators as well (overall market and subsector charts, commodity prices, interest rates, etc.). As suggested in Edwards and Magee (1966), these clues should help the analyst choose the stocks which would be good candidates for trading. Stocks which, from visual inspection of the charts, seem to trend well and make large percentage moves are good candidates. On the other hand, stocks which appear trendless or random, or trade within too narrow a price range, would be less good candidates.

What emerges is a variety of rules and theories, each aimed at describing a certain aspect of the supply and demand interaction. Some rules or patterns may seem to conform to certain price series better than others. Patterns and trends, once identified, do not imply guaranteed price moves - rather, they imply an increased *probability* of that

price move. Some filter rule, formal or informal, should be applied in order to match a trading rule to potential trading candidates. Trade signals should be tempered with clues that can be gleaned from any number of sources. Clearly technical analysis, as prescribed by its practitioners, *is* an art. It defies easy or mechanical implementation. As such, it defies easy or mechanical verification. Thus, as shall be discussed in a later chapter, a formal study of technical analysis presents a formidable challenge.

2.2.3 Theory

As mentioned above, the theory of technical analysis consists in fact of numerous trading rules and rules of thumb. While it would be nearly impossible to cover *all* of these rules, most fall into one of three general categories - trends, support and resistance and patterns. Trends occur due to an excess in supply or demand, and trading these simply involves buying low and selling high. Support and resistance levels emerge when prices seem to consistently bounce off a certain level. These levels tend to be round numbers, leading to the interpretation that investors (as well as perhaps analysts) set pre-defined targets for stocks - levels at which they believe the stock is a buy or a sell, and which to some extent become self-fulfilling. Finally, patterns involve the consolidation of prices due to a relative balance of supply and demand, and together with volume clues, can signal a continuation or a reversal of trend. Each of these three categories of rules will be examined below.

A. Trend Following

The concept of trend following is the most basic element of technical analysis. If it is believed that asset prices do indeed trend, then the primary concern simply becomes

one of identifying trend as early as possible. As outlined above, the Dow Theory and the principle of confirmation form the basis for identifying the trend of the overall market.

As markets have become more sophisticated, technical analysts have modified and expanded on some of the earlier theories. It has become accepted that within an overall market trend, certain subsectors of that market will trend differently - thus the importance of the wide array of indices and averages.

The first *rule* to be adhered to, in terms of following a trend, involves only taking positions that are consistent with the overall trend. That is, if the overall trend is upwards (bull market), one should only take long positions (i.e. buy the security). Likewise, if the overall trend is downward (bear market), one should only take short positions (i.e. sell the security). A common saying is *the trend is your friend* - the implication is simply that you will be unsuccessful if you take positions which are against the trend.

The next consideration involves determining trend within a given price series. The most common means of doing this involves simply drawing a trendline on a chart of past prices. Charts come in numerous varieties - equivolume, candlestick, point and figure, etc. - but the most common is simply the bar chart. A daily bar chart will record the open, high, low and closing prices for the security. As well, the volume is usually recorded below the prices along the same time axis. A trendline runs through closing prices and will connect subsequent troughs if the trend is upwards, or subsequent peaks if a trend is downwards. Figure 2.1 shows an example.

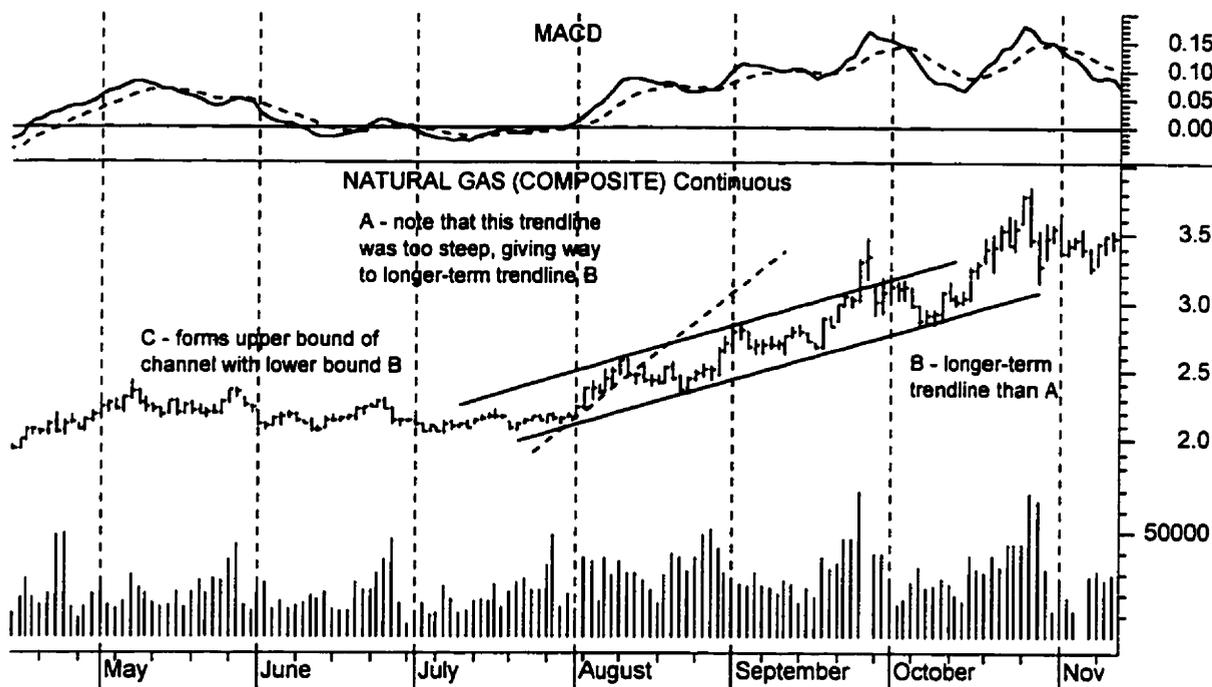


FIGURE 2.1 – Trendlines and Channels

Several points are worth noting in the chart above. The choice of trendline will evolve over time and will depend on the time horizon of interest. Shorter horizon trends will be steeper, but if broken, do not necessarily imply that the trend direction has changed. Note in the above example that the dotted trend A gave way to the solid trend B - also, trendline B proved to hold over a longer period of time. Another point is that a trendline becomes more significant the more times it is confirmed.

The next consideration involves the reversal of trend. Quite simply, this arises when prices violate a trendline. The more important the trendline, the greater the implications of the break. In practice, some filter rule is usually used in order to confirm the breaking of a trendline. One common rule is the *two-day rule* which holds that a break in a trend must hold for two successive days. Other filters involve percentage moves, in which a certain percentage price move is required before a trend is deemed to

be broken. Many simple trading rules derive from this simple concept of trend. The most basic is simply to close out one's position once the trend has been broken, and to reverse one's position once it is confirmed that the trend has been reversed.

A more common variation involves the moving average (MA). The MA can be calculated in many ways - simple, time-weighted, volume-weighted, etc. The MA will be a *smoother* series than the original prices, and will be smoother if longer time periods are used in calculating the average. Clearly, prices will fluctuate above and below the given MA. Another simple trading rule involves buying as prices rise through the MA and selling when prices move below it. A rule which involves less *whipsaws* (i.e. false moves which can generate costly commissions) uses two MA's of differing length. A buy signal arises when the fast MA crosses up through the slow MA, and vice versa.

One of the most popular technical indicators is the moving average convergence-divergence indicator (MACD). The MACD simply plots the convergence and divergence of two pre-specified moving averages. The MACD crosses into positive territory when the fast MA rises above the slow MA, and vice versa. By plotting a moving average of the MACD itself, we can get an earlier indication of which way the prices are moving. See the top panel of Figure 2.1 for an example.

Another concept related to trend is that of channel. It is felt that a given trend will often be contained within a channel, meaning that prices will fluctuate between a lower and upper trendline. Trading a channel is fairly straightforward - buy a bounce off the lower trendline, sell a bounce off the upper trendline. Note that trendlines B and C in Figure 2.1 constitute a channel.

The above summary only skims the surface of the literature, highlighting the most basic of trend following rules. Several points are worth noting. While the concept of trend is so fundamental to technical analysis, it is by definition very ambiguous in terms of implementation. On any given chart, there are a multitude of potential trendlines - the choice of which ones are most significant is subject to the discretion of the particular analyst. Likewise, the reversal of trend is equally subjective. Rules involving MA's give more definitive signals, though with these rules the analyst must decide which lengths of MA are most relevant. Clearly, the implementation of even the most basic of technical trading rules involves a great degree of subjectivity on the part of the analyst.

B. Support and Resistance Levels

Edwards and Magee (1966, p.211) define the concepts of support and resistance:

... we may define *support* as buying, actual or potential, sufficient in volume to halt a downtrend in prices for an appreciable period. *Resistance* is the antithesis of support; it is selling, actual or potential, sufficient in volume to satisfy all bids and hence stop prices from going higher for a time...

Essentially, support and resistance levels are identified on a chart in the same way as are trends. A local maximum (or minimum) occurs, and then may reoccur any number of times. The interpretation is that, for whatever reason, there is a substantial amount of selling (or buying) at that level.

Technical analysts offer several explanations as to why support and resistance levels may exist. The most common involves the investors' mentality or psychology. It

is postulated that investors make a commitment to a stock at a given price, and purchase it. Suppose the price subsequently declines, they are then somewhat unwilling to admit they were wrong. They will very likely wait, hoping to sell at the price he paid - hoping to *get out even*. While this explanation is certainly consistent with observable investor behavior, it does not explain very well why *particular* levels appear as support or resistance.

However, the story becomes more complete when one considers that investors often act in unison - the behavioral finance literature, outlined in Section 2.4, makes this point. In addition, when one considers large order institutional trades and the tendency towards converging analyst expectations, it is easy to see how particular levels could emerge as either support or resistance. For example if several analysts, with institutional buying behind them, recommend IBM at \$100, it is likely that if the price approaches that level there may be considerable buying, perhaps enough to stem a previous decline. Note that the existence of these levels is not inconsistent with rational buying, but may simply be a result of cost efficiency in the information gathering process.

The concept of support and resistance implies that these levels should be *predictable* to a certain extent. For example, if a chart seems to show an important support level existing in the past - say IBM at \$100 - then following an increase to say \$150 and a subsequent decline, one would expect continued support at the \$100 level. If an investor were short the stock, he may wish to cover his short slightly above the \$100 level, feeling that support should arise there, causing excess demand and pushing the price back upwards.

Technical analysts also argue that a previous resistance level, once broken, can reverse roles and become a support level. This phenomenon is explained using the supply and demand framework. Resistance levels are formed when a local maximum is reached repeatedly. Presumably at the resistance level there is a considerable turnover of shares, so as to cause a reversal of price movement, with a large seller but equally a large number of buyers. Once the supply has been bought up, the resistance level is overcome, and those numerous buyers now consider stock at that level to be a bargain such that it becomes a support level. Again, the explanations behind the movements are not that important to the technical analyst - the extent to which any given chart phenomenon can aid in his forecasting is all that matters to him.

An important signal is given when a support or resistance *failure* is observed. That is, when prices move through an important support or resistance level, it may signal a major price move. A most important clue is given by the corresponding volume activity. Essentially, for a price move to be genuine, it must be accompanied by an increase in volume. This is said to be especially true for upward movements in prices - one saying has it that *stocks must rise on volume but can fall of their own weight*. A behavioral explanation of this sort of occurrence would be that investors are happy and quick to take profits, but unwilling and slow to accept losses.

Support and resistance levels are important tools for the technical analyst. On their own they can offer important clues as to prices at which to buy or sell. A break through one of these levels, with corresponding volume increases, can signal a major move. More generally, however, the concepts discussed above play an important role in the formation of price patterns, which is the subject of the following section. Note that,

as with trend, support and resistance can be somewhat ambiguous concepts. In deciding which peaks or troughs define important levels, much is left to the discretion of the particular analyst. The implication, again, is that these are concepts which do not lend themselves well to rigorous testing.

C. Price Patterns

Price patterns are essentially a combination of the forces that are said to cause trend and support and resistance levels. It is believed that when some imbalance between supply and demand exists, the physical turnover of shares between buyers and sellers will result in the formation of recurring patterns. Together with the information provided by volume levels, it is felt that near-term future price moves can be predicted with reasonable accuracy based on these patterns. While there are numerous patterns discussed in the technical analysis literature, most can be classified as either *continuation* or *reversal* patterns.

Due to the focus of technical analysts on the concept of trend, the prediction of trend reversal becomes extremely important. There are a number of clues that may appear on a chart which may hint at an impending reversal of trend. One is simply a lack of confirmation in volume - should a stock make new highs, yet not make new highs in volume, the integrity of the trend may become suspect. More reliable, from the point of view of the analyst, are patterns which reveal a shift in the fundamental imbalance of supply and demand. Together with the belief that history repeats itself - that the same dynamics of trend reversal will occur time and time again - the existence of these patterns becomes a valuable forecasting tool.

Perhaps the most common reversal pattern discussed in the literature is the *head-and-shoulders* formation. This pattern is also representative of many other patterns, in that it incorporates the concepts of trend and support and resistance, as well as the importance of volume, within the concept of pattern. See Figure 2.2 below for an example of the *head-and-shoulders top* formation.

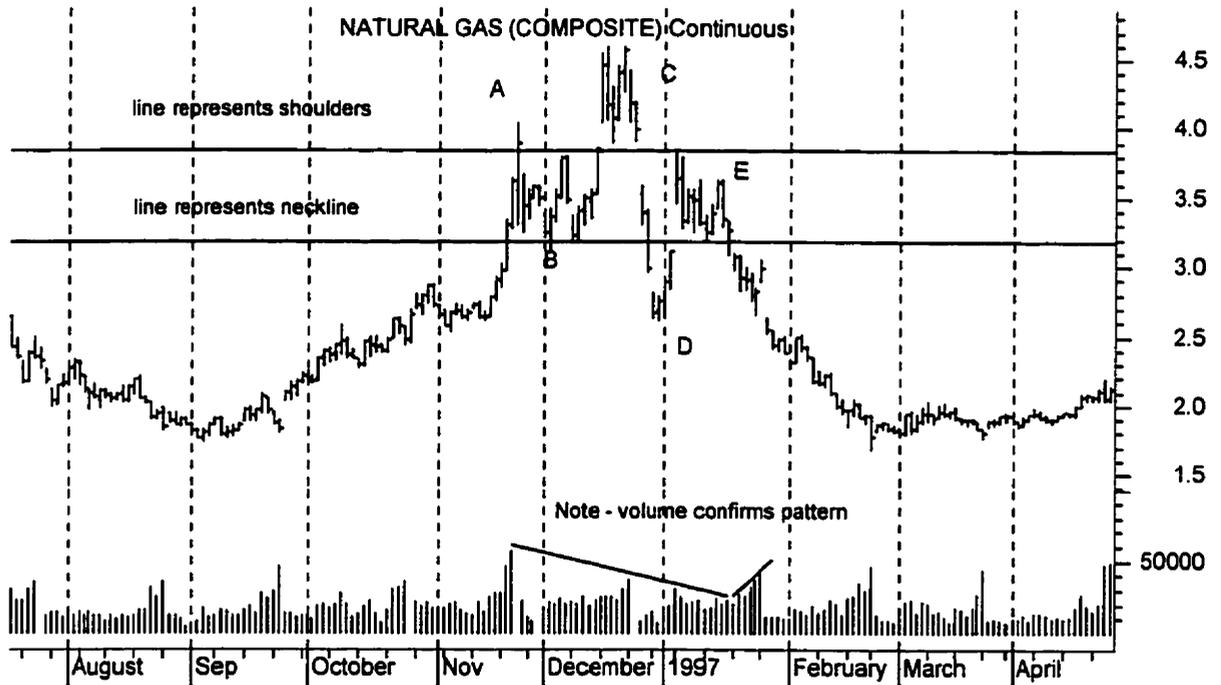


FIGURE 2.2 – Head-and-Shoulders Reversal Pattern

Essentially, the pattern is based on the *lack of confirmation* principle. The rally to peak A occurs on a certain amount of volume. The correction brings prices to level B. At this point, there is no reason to suspect anything. Next, prices rally to a new high at level C, but volumes are lower than at the previous peak at A. This lack of confirmation should put the analyst on the alert - essentially, he would like to see a confirmation of the upward trend. Prices drop only to level D, which confirms this level as support, since

prices bounced there at point B as well. When prices rally again, but only to point E (which is at the same price level as A and therefore confirms a level of overhead resistance), *and* do so on still lower volume than the previous peak at C, the analyst concludes that the previous uptrend has ended. When prices break down through the *neckline* (the support level common to points B and D), and do so on *increased* volume, the analyst concludes that a reversal of trend has taken place.

The interpretation of the pattern is fairly straightforward using the logic applied in the discussion of trend and support and resistance. Essentially, the picture is of an upward trend losing momentum - equally; it is a picture of excess demand being met and then thinning out. Excess demand eventually meets large levels of supply at point A, turning prices back. At point B, with prices lower again, excess demand takes over again, but with lower volume this hints at the fact that excess demand may be waning. This pattern repeats through points C and D. At point E again there is excess supply, and perhaps increasingly so, as sellers begin to feel that previous selling opportunities at A, C or E may have been missed. The break down through the neckline confirms that there is now an excess supply, and that the trend has been reversed from upwards to downwards.

It is this sort of logic - the interplay between buyers and sellers, and between excess demand and excess supply - which characterizes the majority of the patterns in technical analysis. Patterns such as the *double-top* and *triple-top* are simply variations of the head-and-shoulders, but are said to occur with less frequency than the latter. Other patterns, such as the *saucer-bottom*, are again variations on the same basic logic.

Likewise, these patterns are said to exist at both market tops and bottoms, with the appropriate adjustments made to the arguments - i.e. volume confirms the *new* trend, etc.

In addition to patterns showing trend reversal, there are said to exist *continuation patterns*. Essentially, these patterns involve *consolidation* (sideways chart movement) following an earlier move. The logic here is again straightforward. Suppose prices have risen for several weeks due to an excess demand for shares. Presumably some investors will have shorter time-horizons than others, or will have set more modest profit goals for themselves. An increase in supply may arise temporarily - often referred to as profit taking - such that the upward trend is halted. Following this so-called consolidation, the excess demand continues to push prices higher. The key to knowing the different implications of the patterns, argues the technical analyst, lies in the type of pattern and in the volume clues.

Perhaps the most common of the continuation patterns is the *flag* or *pennant* formation. These tend to occur following a sharp price increase. Following a sharp increase is a period of profit taking. There is still high demand, but the profit taking involves temporary excess supply, and prices trade in a tight range declining somewhat. Visually, this appears on the chart as the flagpole and flag. Next, there occurs a sharp increase in prices, similar in nature to the previous increase. Therefore, the formation of a flag or pennant can provide the analyst with an important clue as to an important price move to come.

The key to this pattern, as with the others, lies in the corresponding trend in volume. Patterns are essentially a consolidation, involving a relative balance of supply and demand. As such, volume tends to drop off as the pattern nears completion. There is then typically a dramatic increase in volume, as a new imbalance emerges. Far from being some mystical sequence of prices which the analyst *hopes* will cause a future price

move, patterns are simply the visual record of past movements of supply and demand. The analyst does not attach any symbolic importance to the patterns, but simply feels that from past experience certain patterns may, with increased probability, precede a particular price move.

2.2.4 Summary

The preceding sections have only touched upon the vast body of theory known collectively as technical analysis. There are innumerable trend-following systems which have not been touched upon. As well, only a few patterns have been outlined while many others were not - triangles, rectangles, diamond formations, etc. It is hoped that, by working through the concepts underlying the most basic elements of technical analysis, a general flavor of the study and its logical foundations has been given.

An important distinction that should be made is that pointed out by Murphy (1996) between the *technician* and the *chartist*. Essentially, this distinction involves the sophistication of the tools employed by the analyst. Whereas in its early days, technical analysis had no tools *other* than charts, this is no longer the case. The theories of the chartist, then, correspond to the most basic and earliest theories - ones which can be examined visually on the chart. The theories of the technician, by contrast, are considerably more complex, taking advantage of modern computing capabilities - these include the creation of complex trading systems based on artificial intelligence and neural nets, with the possibility of screening thousands of securities so as to find compatible trading rules. Although the present summary has outlined only the most basic chart

concepts, it should be noted that the fundamental logic behind all technical analysis is quite similar in nature.

2.3 Efficient Markets Theory

The efficient markets theory (EMT) refers collectively to a body of theory aimed at proving the informational efficiency of capital markets. Within the broad heading of EMT are several asset pricing models, as well as several classes of informational efficiency. Under certain assumptions, EMT essentially can be simplified to the random walk hypothesis - thus the common use of these terms as interchangeable. This section will outline the origins and philosophy underlying EMT. The random walk hypothesis, in its various forms, will be examined. Finally, the implications for technical analysis are explored.

2.3.1 Origins

Several sources cite Bachelier (1900) and Cowles (1933) as pioneering the theoretical and empirical foundations of EMT, respectively. On the theoretical side, models of asset prices find their origin in earlier mathematical theories of gambling and games of chance. Thus, Campbell et al. (1997) cite the origins of the martingale model as far back as the Italian mathematician Girolamo Cardano in 1565. While Bachelier was the first to make the theoretical link between these pricing models and market efficiency, the theme only began to make an impact in modern economics with the work of Samuelson (1965) and Mandelbrot (1967).

Fama (1970, p.389) notes that EMT essentially grew out of earlier empirical work - that “...until the Mandelbrot-Samuelson models appeared, there existed a large body of empirical results in search of a rigorous theory”. Studies in the 1950’s and 1960’s, such as those of Kendall (1953), Roberts (1959), Alexander (1961) and Cootner (1964), to name only a few, suggested that the time-series of security prices were generally consistent with a random walk. It should be noted, however, that these studies started out with the assumption that speculation was a *fair game* with expected profits being zero. Clearly, there was a joint-hypothesis embedded within the methodology that provided some of the earliest and most powerful support of EMT, and this fact will have important implications for the way in which these results are to be interpreted.

2.3.2 Philosophy

Fama (1970, 1991) provides the most thorough reviews of modern EMT, stating that “...a market in which prices always “fully reflect” available information is called “efficient”...” (1970, p.383). This definition is concise yet ambiguous, though Fama does go on to expand on the concepts “fully reflect” and “efficient”. Essentially, it is argued that a model of equilibrium prices is needed, in order that prices might be said to fully reflect all information and efficiently reach their equilibrium levels.

Campbell et al. (1997, p.20) prefer the following definition, offered by Malkiel (1992):

A capital market is said to be efficient if it fully and correctly reflects all relevant information in determining security prices. Formally, the market is said to be efficient

with respect to some information set... if security prices would be unaffected by revealing that information to all participants. Moreover, efficiency with respect to an information set ... implies that it is impossible to make economic profits by trading on the basis of [that information set].

This definition is preferred in that it explicitly states that “efficiency”, in this context, refers to informational efficiency, as opposed to Pareto or allocational efficiency. Also, it implicitly suggests that market efficiency can be tested in one of two ways – either by a test based on revealing information, which would be implausible due to practical considerations, or by testing trading systems for economic profits, which would certainly be feasible.

Two methods of testing trading systems have been developed – the first based on an analysis of the track record of *actual* traders or investment managers, and the second based on an analysis of the returns generated by a *hypothetical* trading system. The former would be preferable to the extent that it deals with real-life trading activity, but suffers from the problem that it is difficult to know what information set actually faces the trader, or what set of rules he uses to interpret that information. The latter method is therefore more feasible in practice.

The next decision regards which information set is to be used. Roberts (1967), as outlined in Campbell et al. (1997, p.22), proposes the following definitions relating a particular subset of information with a particular class of efficiency:

Weak-form efficiency: The information set includes only the history of prices or returns themselves.

Semistrong-form efficiency: The information set includes all information known to all market participants (*publicly available* information).

Strong-form efficiency: The information set includes all information known to any market participant (*private* information).

In addition to specifying a particular information set, a model of normal returns or equilibrium prices must be specified. The primary challenge here is to adjust for risk in some way. The benchmark in early tests was the capital asset pricing model (CAPM), attributed to Sharpe (1964), Lintner (1965) and Black (1972). More recent models, some of which have arisen in an attempt to address some of the failures of the CAPM, are the arbitrage pricing theory (APT) of Ross (1965), and the consumption-based CAPM models (CCAPM) of Rubinstein (1976), Lucas (1978) and Breeden (1979). An important distinction is whether expected returns are specified to be constant through time or time-varying. As will be discussed in the next chapter, empirical evidence against the constant expected returns models has led to the growth of numerous time-varying expected returns models.

Once normal returns have been defined, abnormal returns are simply the difference between actual returns and the normal expected returns. To the extent that abnormal returns (i.e. any deviation from normal expected returns) are unforecastable or *random*, market efficiency holds. This notion leads to the *random walk hypothesis*,

which essentially goes hand in hand with EMT. The apparent contradiction between an efficient market and random prices moves is dealt with succinctly in Campbell et al. (1997). Basically, the explanation lies in an understanding of the *law of iterated expectations*. Suppose today's prices P_t are based on today's expectation of intrinsic value V^* , conditional on today's information set I_t :

$$P_t = E[V^* | I_t] = E_t V^*$$

Written one period ahead:

$$P_{t+1} = E[V^* | I_{t+1}] = E_{t+1} V^*$$

Note, however, that *at time t* the set I_{t+1} is identical to the set I_t , by the definition of I_t .

That is, while at time $t+1$ the set I_{t+1} is greater than the set I_t , at time t they are identical.

Therefore, the following is true:

$$E_t[P_{t+1} - P_t] = E_t[E_{t+1}[V^*] - E_t[V^*]] = 0$$

This implies, of course, that there is no *expected* price change, and that any price change should be random. In other words, price changes should be unforecastable. Herein lies the direct challenge to the claims of technical analysis - EMT implies that prices simply are not forecastable.

2.3.3 Random Walk Hypotheses

The previous section has shown that when some model of equilibrium prices is assumed, the EMT can be tested under the random walk hypothesis. Campbell et al. (1997) further refine the concept of the random walk, distinguishing between the type of dependence that can exist between an asset's returns at two points in time. This method

allows us to distinguish between a martingale and three different types of random walk, and can be summarized as follows (from Campbell et al. (1997), p.29):

Table 2.1. Classification of Random Walk and Martingale Hypotheses

$\text{Cov}[f(r_t), g(r_{t+k})] = 0$	$g(r_{t+k}), \forall g(\bullet)$ linear	$g(r_{t+k}), \forall g(\bullet)$
$f(r_t), \forall f(\bullet)$ linear	RW3: $\text{Proj}[r_{t+k} r_t] = \mu$	N/A
$f(r_t), \forall f(\bullet)$	Martingale: $E[r_{t+k} r_t] = \mu$	RW1 and RW2: $\text{pdf}(r_{t+k} r_t) = \text{pdf}(r_{t+k})$

Table 2.1 shows that under all of the above hypotheses, returns are uncorrelated - thus $\text{Cov}[f(r_t), g(r_{t+k})] = 0$. If the functions which generate returns $f(\bullet)$ and $g(\bullet)$ are both assumed to be linear, we have RW3. If only $g(\bullet)$ is assumed to be linear, we have the martingale. If both $f(\bullet)$ and $g(\bullet)$ are allowed to be non-linear, we have the stronger versions of the random walk RW1 and RW2. Note that $\text{Proj}[r_{t+k} | r_t]$ refers to the projection of r_{t+k} onto r_t - the fact that it equals μ means that the only change in r over the time t to $t+k$ is the expected price change, the drift term μ . Also, $\text{pdf}(r_{t+k} | r_t) = \text{pdf}(r_{t+k})$ simply implies that the partial density function of r_{t+k} is the same conditional *and* unconditional on return r_t - that is, future returns are independent of past returns. Each of these hypotheses will be outlined below.

A martingale is essentially a stochastic process such that:

$$E[P_{t+1} | P_t, P_{t-1}, \dots] = P_t$$

Simply put, this is the model of a fair game. There is no expectation of future returns even given complete knowledge of past outcomes. To the extent that it fails to account for the relationship between risk and return, it has been widely considered inappropriate in modeling asset returns. It should be noted, however, that asset returns *are* said to follow a martingale process, though here a martingale hypothesis issue.

Using the Campbell et al. (1997) taxonomy, the Random Walk (RW1) describes the evolution of prices as follows:

$$P_t = \mu + P_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim \text{IID}(0, \sigma^2)$$

where μ is the *expected* price change, and ε_t is random, independent, and normally distributed with mean zero and variance σ^2 . Given the fact that, under investor limited liability, returns will *not* be normally distributed, the RW1 can be rewritten in logarithmic form. Even with this modification, the assumption of identically distributed increments is simply violated by changes - economic, social, regulatory or otherwise - that have occurred in their history. Despite this fact, there *are* numerous techniques for testing RW1 - the Spearman rank correlation test and the Kendall τ correlation test.

The Random Walk 2: Independent Increments (RW2)

is a more realistic description of asset prices. The RW2 differs from RW1 in that the increments are independent but not identical. While this better corresponds to the empirical evidence, Campbell et al. (1997, p.41) point out that if "we place no restrictions on how the increments of the data can vary through time, it becomes virtually impossible to distinguish between the RW1 and RW2 models."

inference since the sampling distributions of even the most elementary statistics cannot be derived.” Filter rules and technical analysis are cited as the most common tests of RW2.

Finally, under Random Walk 3: Uncorrelated Increments (RW3) it is specified that increments are dependent but uncorrelated. This is the weakest of the three types of random walk. An example of a process consistent with RW3, but not RW1 or RW2, would be one in which increments are uncorrelated, but squared increments are correlated. Tests of RW3 are numerous, including autocorrelation coefficients and variance ratios.

2.3.4 Implications

To the extent that the empirical results confirm the random walk hypothesis - and the evidence seems overwhelming, as we shall see in the following chapter - those same results seem to refute the claims of technical analysis. There are certain theoretical considerations, however, which must be kept in mind when interpreting the results.

Most importantly, there is the so-called *joint hypothesis problem*. Fama (1970, p.384) explores this issue, stating that “some such assumption [of expected return model] is the unavoidable price one must pay to give the theory of efficient markets empirical content”, while in Fama (1991, p.1575) it is stated more explicitly that “market efficiency per se is not testable”. Campbell et al. (1997, p.24) state that “this *joint hypothesis* problem means that market efficiency as such can never be rejected”. In short, a model of equilibrium prices or expected returns is needed in order that the efficiency of prices in reaching their equilibrium values may be tested. Any rejection of efficiency can be

interpreted as market inefficiency, an incorrect equilibrium model, or both. Thus, when *minor* anomalies are reported in the results, they are usually interpreted as insignificant departures from the random walk. Another interpretation, such as that in Shiller (1989), is that these anomalies in fact may reveal a significant departure from the usual rational expectations equilibrium asset pricing models - this line of reasoning will be examined in greater detail in the following section.

Another problem lies within the assumptions of the various random walk hypotheses. Consider RW3 - referring to Table 2.1 we see that RW3 holds only if $f(\bullet)$ and $g(\bullet)$ are restricted to being linear functions. Now $f(\bullet)$ and $g(\bullet)$ are the functions which generate the asset's returns at times t and $t+k$. Therefore, under RW3 we *assume* that returns will be linear. Yet, a frequent *conclusion* of these results is that trading rules cannot outperform buy-and-hold. Clearly, trading rules such as those specified in technical analysis are predicated on the notion that returns are non-linear, and that results of tests which *assume* linearity cannot be used to refute rules based on non-linearity. Indeed, once linearity is assumed, the result that buy-and-hold outperforms all other strategies is simply trivial.

If we allow $f(\bullet)$ and $g(\bullet)$ to be non-linear, we must test either RW1 or RW2. It has been argued already that RW1 is implausible from a theoretical standpoint, and that RW2 is nearly impossible to test due to the lack of assumed distributions. Therefore, it is not clear that the random walk hypotheses can effectively be used to refute technical analysis. The most effective refutations of technical analysis will have to come from the testing of properly specified technical trading rules.

2.4 Behavioural Finance

Behavioural finance approaches the issue of asset prices from an entirely different perspective than does EMT. Consider the first sentence in Shiller (1989, p.6) - "Investing in speculative assets is a social activity". From an EMT viewpoint, this may well be true, but is entirely irrelevant. From a behavioural finance viewpoint, however, that statement *is* a logical starting point. The discipline draws on the theories of psychology and social psychology, in attempt to describe the decision-making process of investors. It tries to reconcile the evidence in support of efficient markets, with the commonly observable fact that investors *are not* always making rational decisions. The weaknesses of the econometric methods used to provide evidence for EMT are highlighted, and it is argued that many of the empirical results would *also* be consistent with alternative models of price formation. Behavioural finance by no means sets out to refute EMT in its entirety, but it does set out to rewrite the micro-foundations of EMT, as well as to reinterpret some of its results.

2.4.1 Origins and Philosophy of Behavioural Finance

To a certain extent, it could be said that the origins of behavioural finance date back at least hundreds of years, if not to the origins of speculative investing itself. Popular models of investment have always given a prominent role to the forces of irrational buying and selling, as well as to the aspects of fad and fashion regarding investments. In what are often considered to be two of the classic texts on investing -

McKay's (1841) *Extraordinary Popular Delusions* and de la Vega's (1850) *Confusion de Confusiones* - historical episodes of speculative bubbles are recounted. From John Law's Mississippi Land Scheme to the South Sea Bubbles to the Dutch Tulipomania, investor behavior is described as anything but rational. Even present day media coverage of market activity brims with claims of "panic buying" and an "overvalued market", not to mention "irrational exuberance".

More formal origins of behavioral finance are to be found in the fields of psychology and social psychology. Shiller (1989) reviews much of the early work in these fields, as it applies to investor behavior. Studies of irrational decisions, gambling behavior and group polarization of attitudes all point to the conclusion that otherwise rational people can be led towards irrational behavior under certain circumstances. Amos Tversky, Nobel winner in psychology and often considered a founder of behavioral finance, offers the *prospect theory* in Kahneman and Tversky (1979), in which individuals are likely to view losses and gains very differently, to the point that it will affect their risk-taking behavior.

More recently, work in behavioral finance has moved into the realms of finance and economics. Using the theories from psychology and other disciplines, an attempt has been made to redefine the micro-foundations of economics in a way that is consistent with these alternative models of decision-making. The primary approach has been to reinterpret the *anomalies* uncovered in the EMT research, as well as to examine the statistical basis of that same work. Thus, much of the recent work involves the excess volatility in markets, the alleged overreaction to news announcements, as well as

providing alternative models of bubble-type price movements that are still consistent with the empirical results we have.

2.4.2 Implications

The theory of behavioral finance raises a number of interesting questions, and points out many potential weaknesses of EMT. To the extent that it is based on results from psychology, and what we know of *actual* individual decision-making behavior, it is an intuitively appealing theory. Its greatest weakness, however, is that it presents few testable implications. While it *does* attempt to point out flaws in traditional EMT, it is unable to link those flaws to the theory of behavioral finance in any way which is testable.

In terms of technical analysis, behavioral finance seems to remain silent - perhaps wisely so, as a discipline attempting to gain a foothold in the academic community. However, the similarities between the two theories are striking in a number of ways. The heterogeneous nature of investors, and their sometimes irrational decisions, are common to both theories. The possibility that this irrationality can sometimes cause speculative bubbles is certainly common to both as well. Theories such as Tversky's *prospect theory* are consistent with the support and resistance levels of technical analysis, while theories of *group polarization of attitudes* would explain the price and volume breakout concept.

Unfortunately, the similarities between behavioral finance and technical analysis include the fact that neither are very well defined from a statistical standpoint, and neither

lend themselves well to rigorous testing. Perhaps the most fruitful work to be done in behavioral finance in the near future will be related to uncovering further anomalies of EMT, as well as exploring the potential statistical biases inherent in EMT.

2.5 Conclusion

This chapter has provided a brief overview of several competing theories of asset pricing and market activity. Each theory has differing views as to the dynamics of investor behavior and as to the predictability of asset prices. Technical analysis implies that prices trend due to imbalances in supply and demand, and that future prices may be predictable through the use of a variety of trading rules. Efficient markets theory argues that markets are efficient, that current prices accurately reflect all available information including expectations as to future valuations, and that therefore future prices are unpredictable and random. Behavioral finance seeks to redefine the investor's decision-making process, refuting rational expectations, yet agreeing that future price moves are indeed unpredictable.

These three theories have evolved in considerably divergent manners. As such, it is often difficult to compare the empirical results that each discipline claims supports its own theories. EMT has certainly adopted the most rigorous statistical approach of the three, though recent work shows that some of its statistical basis may be flawed. Behavioral finance works within the same statistical approach as EMT, but interprets the results quite differently. Technical analysis has developed its own set of concepts, definitions and tools which bear almost no resemblance to those used in the academic

community. As Campbell et al. (1997, p.43) argue, “perhaps some of the prejudice against technical analysis can be attributed to semantics.”

The main goal of this chapter has been to evaluate the virtual dismissal of technical analysis by the academic community. It has been argued that the statistical basis of EMT may in some ways be flawed, and that a relaxation of some of its more restrictive assumptions may alter the results significantly. It has also been argued that much of the theory of behavioral finance can be accommodated within the concepts of technical analysis. The conclusion, therefore, is not that technical analysis works or should be used, but simply that on theoretical grounds it should not be dismissed outright, and that perhaps its implications deserve further study.

CHAPTER 3 – RECENT EMPIRICAL EVIDENCE

3.1 Introduction

This chapter is arranged as a survey of recent literature, and as a summary of recent empirical evidence, as pertains to the use of technical analysis in energy futures markets. Unfortunately, as with the previous chapter on background theory, surprisingly little work has been done on *specifically* this topic. While this in some ways makes the task at hand more challenging, at the same time it highlights the need for a study of this sort.

Therefore, again as with the previous chapter, the approach must be somewhat piecemeal. For example, the majority of the work on technical analysis has dealt with either stock prices or exchange rates. Efficient markets theory, often seen as sufficiently well supported by the evidence so as to reject technical analysis outright, has come under attack as a number of its *anomalies* have been further explored. In terms of energy futures, much of the work has focused on market efficiency – unfortunately much of this has looked at efficiency in terms of the futures contract's ability to accurately reflect future spot price, whereas a study of technical analysis looks more at weak-form efficiency using the series of futures prices alone.

The chapter will be arranged as follows. Section 3.2 will review some of the recent work on efficient markets theory. Of most interest in the case at hand are some of the attacks on EMT which focus on certain anomalies – indeed, it is against the backdrop of possible EMT shortcomings that technical analysis must be studied. Section 3.3 looks

at recent studies of technical analysis. For the most part these have been tests of trading rules using stock prices, and are interesting in that they highlight progressively more sophisticated means of testing these trading rules. Section 3.4 surveys the work that has been done in terms of market efficiency in energy futures markets. The final section attempts to tie together the evidence in order to reach some conclusions as to the applicability of technical analysis to futures markets.

The goal of this chapter is analogous to that of the previous chapter. In the previous chapter an attempt was made to show that technical analysis should perhaps not be rejected outright due to *a priori* theoretical considerations and that a healthy debate exists amongst the various theoretical camps. Likewise, the present chapter seeks to show that there is also sufficient debate in terms of the empirical evidence, such that technical analysis should also not be rejected outright based on current evidence. In neither case is the goal to argue in favor of technical analysis, nor to vindicate its theories. The goal is simply to argue that neither theoretical nor empirical considerations imply that technical analysis *does not* warrant further study, given that this often seems to be the approach taken in the academic community.

3.2 Efficient Markets Theory

Recent years have seen the EMT come increasingly under attack. Certainly, there is now a consensus that strong-form efficiency is surely false, although this fact alone does not present a serious challenge to market efficiency *per se*. More importantly, there seems to be a growing consensus that “as literal statements of fact, the weak and semi-

strong forms of the efficient markets hypothesis are also false” (French 1992, p.788).

Indeed, as Fama (1991) argues, market efficiency “is an extreme hypothesis, a point on a continuum, and so almost surely false. The interesting task is not to accept or reject market efficiency, but to measure the extent to which the behavior of returns departs from its predictions”.

The approach currently being taken is to investigate whether excess returns can be earned, after appropriate information gathering and processing costs and transaction costs have been accounted for. What emerges is a view of markets that are sufficiently efficient such that any inefficiencies are not exploitable for excess profits. Campbell et al. (1997) use the analogy of relative efficiency as applied to mechanical systems in physics, where certain frictions and losses of efficiency are to be expected. Overall, this is certainly a less exact and rigorous version of market efficiency than that held in earlier years, but the current interpretation is one which is certainly more realistic, and one which need not change many of the fundamental conclusions of the EMT empirical evidence.

This section follows the methodology laid out in Fama (1991). In his summary of the literature on efficient capital markets, he distinguishes between tests for *return predictability*, results from *event studies* and tests for *private information*. These correspond, respectively, to the earlier categories of *weak*, *semi-strong* and *strong-form efficiency*.

3.2.1 Tests for Return Predictability

Recent years have been witness to a renewed interest in studies of return predictability. For the most part, these studies have focused on return predictability as a function of past returns, as well as variables such as price-earnings ratios and dividend yields. As such, these studies directly address the debate over the validity of both technical and fundamental analysis. To the extent that returns *are* predictable in any significant way, this would constitute evidence *against* the EMT.

A number of studies have looked at return predictability given past returns. These are essentially studies of weak-form efficiency, and get at the heart of the debate over technical analysis. While recent work (see Campbell et al. (1997)) shows that returns tend to be positively autocorrelated, Fama (1991) argues that these results may be partly due to the size-effect as well as nonsynchronous trading effects. Basically, positive autocorrelations tend to be more pronounced in returns on the stocks of smaller firms. The shares of smaller firms trade in lower volumes, and this lack of liquidity can mean that the expected and rational response to certain news may be reflected more slowly in the prices of these shares. Thus, whereas returns are measured using daily closing prices, it may be the case that for less liquid stocks we overestimate the degree of autocorrelation.

In practical terms, this means that while a study showing positive autocorrelations may imply that certain prices trend and can therefore be traded profitably, it may simply be the case that the stock is not liquid. That is, it may not be possible to buy on today's open at the price of yesterday's close, and then sell at today's close – it may simply be

that the *very next* trade after yesterday's close is today's closing trade. In short, evidence on autocorrelation in returns needs to be carefully interpreted.

Campbell et al. (1997) show that variance ratios can also be used as tests of RW3. Essentially the variance of returns at different times should be approximately equal. Therefore the ratio with numerator the sum of variances at times t plus $t-1$, and denominator *twice* the variance at time t , should be statistically equal to one. The authors show that, using this test, the random walk hypothesis is rejected at any level of significance in the cases of indexes and portfolios. Interestingly, they find weak *negative* autocorrelation in the case of individual securities, but this is statistically insignificant in this case meaning they cannot reject RW3. They attribute this to company specific *noise* and other market microstructure factors. Overall, the authors conclude that "recent econometric advances and empirical evidence seem to suggest that financial asset returns are predictable to some degree" (p. 80).

Despite the implications of the size-effect and nonsynchronous trading effect, among others, Fama (1991, p.1580) concludes the "recent research is able to show confidently that daily and weekly returns are predictable from past returns. The work thus rejects the old market efficiency – constant expected returns model on a statistical basis". Fama argues later that these results do not imply that rational expectations or EMT should be thrown out, but rather that perhaps constant expected returns should be replaced with time-varying expected returns. This topic will be discussed in more detail below, though it should be noted that, given the joint-hypothesis problem, allowing for time-varying expected returns will make any rejection of EMT an even more onerous task.

While the early work tended to focus on short-horizon returns, much of the recent work focuses on long-horizon returns. Although recent studies do show short-horizon returns to be positively autocorrelated, it is often argued that these autocorrelations are still close to zero and are therefore economically insignificant. Shiller (1984) and Summers (1986) argue that low autocorrelation in returns could also be consistent with irrational price bubbles, in which prices exhibit large swings away from fundamental values. The crux of these arguments is that irrational and faddish behavior is consistent with what psychology tells us, and that due to the joint-hypothesis problem the statistics could be consistent with models *other* than the EMT.

Fama (1991, p.1581) argues that the Shiller–Summers models lack statistical power – “Even with 115 years (1871-1985) of data, however, the variance tests for long-horizon returns provide weak statistical evidence against the hypothesis that returns have no autocorrelation and prices are random walks”. One must wonder, however, what *other* data can we use? That is, if a data problem means that we cannot get strong statistical evidence *against* a null hypothesis, should that in any way strengthen our support of that null hypothesis? Further, even if we accept swings in prices as a fact, these swings can be attributed to either irrational bubbles or to rational time-varying expected returns. As Fama concludes, “a ubiquitous problem in time-series tests of market efficiency, with no clear solution, is that irrational bubbles in stock prices are indistinguishable from rational time-varying expected returns” (ibid).

Work by DeBondt and Thaler (1985, 1987) also explores the possibility that prices exhibit irrational swings away from fundamental values. Their method is to examine the returns of stocks that have underperformed the market. They find that these

underperformers tend to outperform the market in subsequent years. Is this evidence of irrational bubbles? Critics argue that proper risk-adjustment would account for the alleged excess returns. However, given that the risk-return trade off is itself poorly understood, it seems that this issue is far from settled, if it ever can be.

A number of studies attempt to forecast returns using variables other than past returns. These include variables such as expected inflation, dividend yields, short-term interest rates, etc. (not to mention women's skirt lengths, sunspot activity, Super Bowl winners). While forecasts using price-earnings ratios and dividend yields do seem to have some predictive power, Fama (1991) again questions the statistical power of these results. Fama instead puts forth the idea that time-varying expected returns are consistent with rational expectations and efficient markets *within* an intertemporal asset pricing model, such as the consumption based CAPM of Breeden (1979).

Finally, a very active area of research has centered on certain *anomalies* in returns. Monday returns are lower than other days, while the last day of the month and days before holidays have higher returns. The *January effect* suggests that returns are higher in the first few days of the new year, though critics argue that here again we are seeing the size-effect. Reinganum (1992) provides a thorough summary of these anomalies. Fama (1991) suggests that many of these anomalies may be explained through studies of market microstructure. Further, he hints that certain results found in the literature may be the result of data mining, although this claim is left unsubstantiated.

To summarize, it seems that recent evidence shows that returns *are* predictable to some extent. However, certain impediments exist in terms of interpreting this predictability. A limited data set, combined with an inevitable joint-hypothesis problem, mean that by

and large we cannot reject market efficiency. Campbell et al. (1997, p.80) state that “time-varying expected returns due to changing business conditions can generate predictability. A certain degree of predictability may be necessary to reward investors for bearing certain dynamic risks.” While predictable returns may be rational, under certain assumptions, and while rational expectations are consistent with an efficient market, how can predictable returns be efficient? If returns were truly predictable they would not exist, precisely because markets are efficient.

Perhaps the most fruitful line of research is to further explore market microstructure issues. The market efficiency issue aside, the fact that returns *are* predictable would imply that there may indeed be some role for technical analysis. Perhaps this is because of irrational bubbles, or perhaps it is due to market microstructure – the actual reason for return predictability, while interesting, is somewhat beside the point. The point here is that recent research on return predictability actually suggests there *might* be a role for technical analysis, contrary to what is often assumed – i.e. technical analysis is often discredited precisely because it is assumed that recent evidence suggests that it should be. While a number of recent tests of trading rules have failed to produce profits (though others have succeeded), this is not evidence against technical analysis. If anything, it is evidence that the use of those particular rules in the particular manner prescribed in the given study would not be successful. The fact remains that return predictability opens the door for the application of technical analysis.

3.2.2 Event Studies

Event studies are designed to measure the immediate impact of news or an event on the value of a firm. The inherent assumption is that “given rationality in the marketplace, the effect of an event will be reflected immediately in asset prices” (see Campbell et al. (1997) for a thorough review of the subject). Modern event study analysis follows the methodology as set out in Fama, Fisher, Jensen and Roll (1969) and can be summarized (again, Campbell et al. (1997, p.151)) as having seven steps:

1. Event definition – the event itself and window of time of interest
2. Selection criteria – which data to use
3. Normal returns – need a model to define normal returns
4. Estimation procedure – estimate model out of sample
5. Testing procedure – define hypotheses and statistical tests
6. Empirical results – usual econometric format
7. Interpretation and conclusion – results versus theory

Clearly, we can see that event studies will be laden with the same joint-hypothesis problems as any study of market efficiency. First, we assume rational markets such that price reaction to an event is quick, complete and meaningful. Second, we must assume some equilibrium model of returns that defines normal returns. Surprisingly, Fama

(1991, p.1602) claims that “event studies are the cleanest evidence we have on efficiency (the least encumbered by the joint-hypothesis problem)”. He also claims that “daily data allow precise measurement of the speed of the stock-price response – the central issue for market efficiency. Another powerful advantage of daily data is that they can attenuate or eliminate the joint-hypothesis problem, that market efficiency must be tested jointly with an asset-pricing model” (p.1601). This seems a rather baffling interpretation of market efficiency and of event studies in general. Surely it cannot be the case that *any* rapid response to information serves as a stand-alone proof of market efficiency.

Despite any problems that the joint-hypothesis problem may cause, a number of interesting results have emerged from event-study analysis. Evidence shows that takeover candidates will experience rapid appreciation in share prices. Dividend decrease announcements elicit negative returns (Ahrony and Swary (1980)), counter to the Modigliani-Miller (1969) Irrelevance theorem and counter to rational taxation considerations. If a signalling theory (for example Miller and Rock (1985)) is invoked, however, the reaction of the share price is again rational. This example highlights the joint-hypothesis problem. Theory suggests a given return, but the empirical results show otherwise. We then accommodate the results with a theory that directly counters our original hypothesis, and together theory and results are used to confirm rationality and market efficiency.

While most event studies seem to confirm that prices react rapidly to news (and we assume these reactions to be rapid and correct), a number of studies find excess variance in returns surrounding news as well as post-announcement drift. Not surprisingly, there are a number of theories as to what any anomalies may or may not

mean. A reasonable story is that in certain cases the market *cannot* know precisely what impact an event will have on a firm. For example, a regulatory change may thrust a firm into a more competitive environment. This could affect the firm in a number of ways, and the market can only wait and see how the firm responds over time. Excess variance and price drift could easily be consistent with rationality and market efficiency.

To summarize, the results generally show that the market responds rapidly to news with obvious implications. Unfortunately, event studies cannot shed light on the debate between market efficiency and competing theories. Event studies, however, have been and should continue to be a valuable source of results and information, despite any debates that may arise over the interpretation of those results.

3.2.3 Tests for Private Information

This category corresponds to the earlier category of strong-form efficiency. Fama (1970) points out two cases of private information being used to generate excess profits – corporate insiders with inside information and market specialists with a monopoly on information regarding orders. Neither of these is surprising, and as such were seen as violations of strong-form efficiency, but not as serious threats to EMT in general. Since that time, however, a number of new studies have come to light. Not only do we now have more accurate information on insider trading, but we have new evidence which suggests that money managers and security analysts *may* have access to information that is not public.

In terms of insider trading, Jaffe (1974) finds that indeed corporate insiders can profit from their private information. More importantly, he finds that the market does not respond quickly to *public* information regarding the activity of insiders. His conclusion that outsiders can profit for up to eight months following insider activity, however, is disputed on the grounds that his choice of equilibrium returns model, the Sharpe-Lintner-Black model, is inappropriate – again, the joint-hypothesis problem.

Much work has been done in which the information revealed by the *Value Line Investment Survey* is shown to be inconsistent with public information (see Fama (1991) for a review). The top ranked stocks in this survey *do* seem to outperform the worst. Furthermore, ranking changes produce a significant information effect.

The interpretation here is rather interesting. A major source of information for the investment community announces that a stock is undervalued. The stock then rises in price. Is the price move because the survey was correct, or because the announcement itself generated additional interest in the stock? Is this truly an information effect, or might it more properly be seen as a publicity effect? Furthermore, given that the announcement has the ability to be self-fulfilling, how do we know that *Value Line's* private information was in any way correct?

Finally, a number of studies have looked at professional portfolio management. Results here are very mixed, with some funds outperforming major market indexes while others do not. Further complicating the matter is the factoring in of appropriate costs as well as making the appropriate risk-adjustments, given that variances will vary across funds and indices. As a general statement, however, it would seem that few, if any, funds consistently outperform the market indices.

3.3 Technical Analysis

Studies of technical analysis have also seen renewed interest in recent years, partly as a result of the EMT coming under attack. By and large, the literature on technical analysis has evolved along very different lines than has the academic literature on market efficiency. Campbell et al. (1997) suggest a number of reasons why this may be the case. On the one hand, it may to a certain extent be simply a difference in semantics, in which each body of literature has developed its own set of terms and definitions. On the other hand, it is certainly true that the market efficiency literature has been developed within the more formal and rigorous statistical framework. As Campbell et al. (1997) point out, however, tests of technical trading rules are perhaps the only type of testing which can be used to test RW2 – the case of independent but not identical increments – given that statistical inference in this case is very complex.

This section will review some of the specific work that has been done on technical analysis. In general, this work consists of the testing of specific, simple trading rules. Unfortunately, the results of these studies can only serve as evidence *against* the use of particular rules, but not as a general statement on the validity of technical analysis. Furthermore, it will be shown that most of the tests of specific trading rules have been carried out in ways that are not consistent with the general methodology of technical analysis. Other interesting work has been done in terms of developing more formal testing procedures, as well as attempting to test trading rules in ways more closely in line with the ways prescribed in the theory. What follows, then, is a brief summary of several

of the more important studies on technical analysis that have appeared in the recent literature.

3.3.1 Neftci (1991)

This study aims at formalizing some of the rules of technical analysis. The author argues that a number of these rules are not well defined. In addition, an attempt is made to resolve the paradox that technical analysis is widely used within the financial community, yet much academic work shows it to be inferior to other linear forecasting techniques.

Using the concept of Markov times, the author is able to show that a number of the rules of technical analysis are not well defined. Essentially, “Markov times are random time periods, the value of which can be determined by looking at the current information set (and) cannot depend on future information”. The point being made is that an analysis of charts can pinpoint a minimum or a maximum, whereas these are not Markov times, as future information is needed to *recognize* the given maximum or minimum. What this means is only that *naive* interpretation of charts suggests rules which are not Markov times. As will be recalled from the theory of technical analysis as outlined in an earlier section, the process of generating buy and sell signals is a dynamic one requiring continual reinterpretation (whether subjectively, or objectively through a series of filters with preset parameters) of past price movements and principles of confirmation, but one which is entirely consistent with Markov times.

The author finds that whereas rules involving trendline penetration are not well defined, those involving moving average crossovers *are* well defined. In the absence of complex filtering devices that model the subjectivity of the analyst, it is certainly true that moving average rules are *more easily* defined. The next question involves the usefulness of a technical trading rule that is indeed a Markov time. The author shows that, under the assumption that the observed data can be characterized as a linear process, there will be no use for technical analysis. This is not surprising, however, since *clearly* the basic premise of technical analysis involves the view that asset prices move in a non-linear sequence. What is surprising is that a number of statistical tests aimed at disputing the technical trading rules *assume* linearity, although this is inconsistent with the very notion of active trading.

Using a 150-day moving average crossover rule on the Dow-Jones index, the author finds that this trading rule does indeed have some predictive power. Surprisingly, the author does not elaborate on the implications of his findings. He simply concludes that “if the processes under consideration were nonlinear, the rules of technical analysis might capture some information ignored by Wiener-Kolmogorov prediction theory”. Presumably, if the Dow-Jones index exhibits nonlinearities that can be exploited by technical analysis, the interesting question becomes “is linearity an appropriate assumption in terms of characterizing financial asset prices?” Despite the unanswered questions, the study is valuable in that it highlights the importance of formalizing the technical trading rules, as well as showing that these rules can be of value if the assumption of linearity is relaxed.

3.3.2 Taylor and Allen (1992)

The authors conduct a study of the London foreign exchange market on behalf of the Bank of England. They carry out a survey of over 400 senior foreign exchange dealers, designed to reveal both quantitative and qualitative measures of the importance of technical analysis in this market. The response rate is quite high at over 60%, and of these respondents, 90% respond that they use technical analysis to some degree. The results therefore show a striking reliance on technical analysis, especially over shorter (intraday to one week) forecast horizons. Over longer time horizons, forecasters shift their focus towards economic fundamentals. Other results showed that the dealers saw fundamental and technical analysis as complementary, and that to a certain extent technical analysis was likely self-fulfilling.

The same authors find in an earlier study (1990) that over short horizons of one to four weeks, the forecasts of technical analysts outperformed a number of alternative forecasting methods, including the random walk, vector autoregression and univariate autoregressive moving average time series models. In this same study they find that the actual forecasts vary considerably across analysts, although this is hardly surprising given the vast array of tools and theories employed by the technical analyst. This latter finding casts doubt on the idea that technical analysis may be self-fulfilling – if technical analysts cannot agree amongst themselves, they can hardly be credited with creating self-fulfilling forecasts (as argued in Murphy (1996)).

Perhaps the most interesting result of this study is simply the extent to which technical analysis is currently being used in the financial community. Any basic

institutional approach to economics will lead to the notion that if a service exists and is being paid for, then it *must* be of some value. Arguments of the type put forth by Malkiel (1990) – that chartists are employed in order to generate trades and commissions – are entirely inadequate. The Taylor and Allen survey shows that firms hire chartists because they believe them to be of value in forecasting, and that firms employ these forecasts in trading their *own* assets and not only those of their clients. Clearly, a more plausible story is needed in order to account for the prevalence of technical analysis within the financial community.

3.3.3 Brock, Lakonishok and LeBaron (1992)

The authors test two popular technical trading rules using the Dow-Jones index series. Standard statistical techniques are complemented by the use of bootstrap techniques. The results support the use of these trading rules, and show that the returns earned by these rules are not consistent with a number of commonly used null models of return distribution. The study is interesting both in terms of its conclusions regarding the usefulness of trading rules, and for its innovative methodology.

The study tests two of the more common trading rules – a simple moving average rule and a trading range breakout rule. While traditional tests of these rules show that they are indeed profitable, the authors go a step further. By employing the bootstrap methodology, they are able to “develop a joint test of significance for the set of trading rules”. Essentially, the bootstrap methodology involves “scrambling” the data to form a

new series which can be compared to the original series or can be fitted to a variety of null models. Using this technique, the authors are able to show that the returns earned by the trading rules are inconsistent with four commonly used null models – the random walk, the AR(1), the GARCH-M and the EGARCH. Clearly, this study presents problems both in that it suggests weak-form inefficiency and in that it questions the validity of some of the most common assumptions on return distributions.

3.3.4 Blume, Easley and O'Hara (1994)

The authors examine the role that volume plays in technical analysis. Using a model of aggregate supply and demand, it is argued that volume conveys information not contained within the price statistic – information referred to as the *quality* of the price information. Their theoretical model is consistent with empirical findings of a V-shaped price-volume relation in asset prices. Basically, higher volumes correspond to large price moves. Furthermore, dispersion *decreases* towards these high-volume large price moves, implying that higher volumes signify a greater precision in price information.

These results show why observing price and volume together may be of more use than using price information alone. Interestingly, the authors point out that the role they find for the volume statistic “is remarkably similar to that claimed by proponents of technical analysis”. Indeed, volume information is considered essential to technical analysis, as outlined in Edwards and Magee (1965) and Murphy (1986). It should be noted, however, that the majority of academic studies of technical analysis have *not*

incorporated the volume statistic. Rather than speak of “naïve trading rules”, as in Neftci (1991), we might more properly speak of naïve implementation of trading rules. This study, therefore, highlights not only the importance of volume information, but by implication points out a major weakness in the majority of the studies of technical analysis. It is interesting to note that the conclusion that volume information improves predictions is consistent with the finding in Serletis (1992) that “knowledge of past trading volume improves prediction of futures price volatility beyond predictions that are based on past futures price volatility alone”.

3.3.5 Taylor (1994)

In this study, the author tests the commonly used channel trading rule on currency futures. This trading rule is based on the principle in technical analysis that prices trend, and that they often trend within a *channel* bounded by an upper and lower trendline. The author attempts to show that this trading rule produces excess profits, and that its predictions outperform those made using an ARIMA prediction rule.

The author provides several reasons as to why the channel rule may be profitable, and why it outperforms ARIMA predictions. To quote the authors, “the channel rule may be superior when a trading objective is evaluated because it may require less information to learn about a satisfactory value for its one parameter than an ARMA model needs to find satisfactory estimates of its AR and MA parameter”. An interesting result is that the channel rule can be quite profitable even when prices move in a highly uncorrelated manner – one which *resembles* a random walk. This is because the channel rule does not

make predictions of future prices, but only of the direction or sign of future price changes. If these predicted signs are correct often enough, profits can be made. The author shows that “more than 60% of the sign predictions can be correct for the channel rule even when the maximum correlation between returns on different days is less than 0.02”.

This study is interesting for a number of reasons. The author presents a technical trading rule which, while still quite simple, is nonetheless more sophisticated than many of those used in earlier studies. Also, as a benchmark, the author compares the predictions of technical rule with those of the popular ARIMA models, and finds the former to compare favorably. The author concludes that these results may indicate an inefficient market *or* may be evidence of central bank activity within these markets.

3.3.6 Studies in *Technical Analysis of Stocks and Commodities Magazine*

Technical Analysis of Stocks and Commodities is a trade magazine devoted to the practice of technical analysis, and aimed at practicing traders. It is not an academic journal, and as such presents its material in a very different manner than would an academic publication. Unfortunately, this makes comparisons of ideas between the two types of literature fairly challenging. Nonetheless, the magazine *does* review a number of interesting trading models, and presents (in its own way) serious challenges to proponents of the efficient markets and random walk theories.

In general, the models tested in this magazine are much more complex than those tested in academic journals. This is not to say that they are better, only that they are more closely related to the theory of technical analysis. Recent articles examine the applicability of neural nets (see Rubino and Nimey (1996)) and artificial intelligence. The main advantage with these sorts of systems is that they allow for dynamic reinterpretation of a set of rules within preset (or perhaps evolving) parameters. While they risk becoming overly complex, and at times misleading, they are perhaps the only methodology capable of truly modeling the variety of rules that constitute technical analysis. It will be interesting to see how this literature will evolve, and to what extent the academic research will follow suit in testing these more complex trading systems.

3.4 Energy Futures Markets

There has been increased interest over the past several years in gaining a better understanding of the movements of energy prices. Given the risk of supply shocks and the growth in energy trade, as well as a continual process of regulatory change, this interest is not surprising. As the energy markets evolve, they become more sophisticated through increased liquidity, increased numbers of contracts and the development of an active derivatives market. The literature on energy markets has therefore been extensive, in an attempt to better understand this complex market in the light of these numerous changes.

In terms of market efficiency, considerable work has been done. Most of this work, however, pertains to efficiency in the context of the ability of a futures price to accurately reflect a future spot price. Much less work has been done on market efficiency in the context of weak-form efficiency – that is, the ability to profitably trade a given futures contract based on the past market action of that same contract. This, of course, has been a recurrent theme in this chapter and the last. Simply put, not much work has been done in terms of examining the usefulness of technical analysis in energy futures markets. Nonetheless, a number of interesting studies have been done, and a survey of these can certainly help in gaining a better understanding of the movement of energy futures prices.

This section consists of a survey of selected recent studies aimed at improving our understanding of these prices. The two aspects of futures market efficiency will be touched on, as will certain aspects of trading volume and volatility. Together, it is hoped that these studies will shed some light on the type of work that has been done in terms of characterizing the movement of energy futures prices.

3.4.1 Gjolberg (1985)

The author tests for weak-form efficiency in the Rotterdam spot market for gas oil. Essentially, the hypothesis of efficiency is rejected on the basis that returns are found to be serially dependent. A simulation is carried out, in which a simple trading rule based on runs and reversals (or serial dependence) is shown to generate excess profits.

It should be noted that this study suffers from a number of problems, however. In terms of data, the simulation covered only half of a year of prices, which is hardly enough evidence to prove that a trading rule will be consistently profitable. Furthermore, the profits generated by the trading are somewhat ambiguous. No percentage return is given, no mention of margin or associated costs is made, nor is any comparison with a buy and hold strategy made. Nonetheless, the study does perhaps highlight some inefficiency in the Rotterdam spot market for gas oil. Clearly, it would be interesting to extend this sort of study to include more sophisticated techniques, as well as to generalize it to other energy markets.

3.4.2 Bopp and Sitzer (1988)

This study examines the efficiency question in terms of the informational contribution of a futures contract to its respective spot market. In other words, does a futures market add to the efficiency of a spot market by revealing information about future spot prices? The authors examine this issue by forecasting spot oil prices, with and without the futures price as an explanatory variable. Their general finding is that the futures market does indeed contribute to increased efficiency in the spot market.

The approach taken in this study is interesting in itself. It takes on the efficiency issue from an unusual perspective. No attempt is made to test market efficiency *per se*. Rather, the point of interest is whether a futures market adds to the efficiency of a spot market, and it is found that it does. This work sets the stage for future work examining more closely the relationship between spot and futures prices.

3.4.3 Panas (1990)

The author examines the Rotterdam and Italian oil markets for weak-form efficiency, in an attempt to complement the work of Gjolberg (1985). This study is considerably more sophisticated in terms of its statistical approach than is its predecessor. Tests for serial correlation, unit roots, runs and stability are all carried so as to answer the efficiency question more definitively. In general, the author finds that the data *do* correspond to an efficient market model, and that they correspond to a random walk model.

This study does indeed complement that of Gjolberg. The author examines different oil prices on different exchanges, and gets different results than does Gjolberg. This, of course, leaves a number of questions unanswered. Are the markets for gas oil and other types of oil significantly different? Does the different methodology used in the two studies account in part for the different results? This study certainly contributes to the debate over market efficiency in the markets for energy futures, and is important for that reason.

3.4.4 Serletis (1992)

This study looks at the issue of maturity effects in energy futures contracts. In general, the results support the theory that prices become more volatile and that volume increases as the contracts approach maturity. That the maturity effect does or does not

exist is important for several reasons. One implication of increased volatility concerns the setting of margin requirements, while another concerns the pricing of options. A better understanding of any predictable patterns in volatility would be of considerable interest in addressing these issues.

Having provided evidence for the maturity effect, the study moves on to the issue of causality. The author examines the direction of any possible causality between increased volatility and increased volume. In general, the results show that trading volume Granger-causes price volatility, while the reverse is not true. The important implication here is that past trading volume is established as an important explanatory variable to be used in predicting futures price volatility. This is especially important in terms of the pricing of options. While not a study of market efficiency *per se*, the results do provide evidence of predictable elements within these futures markets. As an extension, it would be interesting to explore what effect, if any, this predictability might have on the futures or other derivatives markets.

3.4.5 Nainar (1993)

This study combines elements of the Gjolberg (1985) and Bopp and Sitzer (1988) studies which preceded it. The author finds that the futures contract *does* add information to the spot market. In addition, spot price volatility increases with trading in futures contracts. Finally, in terms of weak-form efficiency, the author finds that while futures contracts increase volatility and therefore increase the potential for trading profits, these profits would likely not cover transaction costs.

This paper is interesting in that it attempts to combine several aspects of market efficiency, although this same fact is one of the weaknesses of the paper. In particular, the discussion of weak-form efficiency is not satisfactory. Basically, the author tests a very simplistic trading rule, shows that in some cases it *does* generate profits, and then offhandedly dismisses these profits as insignificant. Given the difficulty in testing for market efficiency, as outlined in Fama (1970) among others, this sort of conclusion is not helpful. Nonetheless, the paper does present interesting results.

3.4.6 Herbert (1995)

This paper complements that of Serletis (1992), following a similar approach but looking at natural gas futures prices rather than oil futures. In general, the results of this study confirm those of Serletis. In short, trading volume is shown to cause futures price volatility, rather than vice-versa. While the results of the two studies differ somewhat, this is not surprising given that different commodities are used in each. More importantly, the main conclusions confirm one another. The implications for traders and regulators are important.

It is interesting that this study shows a similarity in the behavior of two different types of energy prices. This issue is addressed by Serletis (1994), in which he concludes that “it is appropriate to model energy futures prices as a cointegrated system”. This is an important contribution in that it adds to our understanding of what moves energy futures prices, and this is certainly pertinent to the issues of predictability and market efficiency.

3.5 Conclusion

This chapter has provided a survey of a variety of literature relating to technical analysis and energy futures markets. As stated previously, there is very little work dealing with *specifically* the role of technical analysis in predicting energy futures prices. As such, the approach has been to review what literature is out there, and then attempt to draw some common conclusions.

The review of recent work on market efficiency leads to one conclusion – that cracks are beginning to show in the EMT, as it is usually known. This is not to say that we should throw out EMT in favor of a competing theory. In fact, EMT holds up very well in most studies, even under its sometimes extreme assumptions. What it means, however, is that we should actively explore the persistent anomalies that seem to exist, and to see what implications these have for market efficiency. In particular, issues relating to market microstructure may shed light on some of the observed anomalies.

There has been a resurgence of interest in the validity of technical analysis. Given that *any* validity for technical analysis would contradict EMT, this recent interest is likely related to the attacks on EMT. A number of studies have shown that even simple trading rules can have predictive power, although the profitability of these rules over time is less well established. Recent work has also focused on testing more elaborate trading rules, which is important given that a number of the earlier studies tested rules that were quite clearly at odds with the theory of technical analysis.

Finally, while there have been few studies of weak-form efficiency in energy futures markets, there have nevertheless been a number of studies which have contributed

to a better understanding of the movement of prices in these markets. We have mixed results as to the efficiency of these markets, in general, although we should remember that a number of these markets are recent and are in a rapid state of evolution. An interesting result concerns the cointegrated nature of energy futures prices, and this could play a role in better understanding what drives these prices.

Combining these results, what can we conclude about the question at hand – can technical analysis be used to predict energy futures prices? The short answer is, of course, that we don't know. What we *do* know is the following. EMT is not without flaws. Numerous studies have pointed out anomalies within the theory, and have pointed out the extremely restrictive assumptions it entails. While EMT is often cited as a reason why technical analysis cannot work, it is clear from recent empirical work that this conclusion may be premature. It is *possible* that technical analysis may capture some anomalies or market microstructure phenomena in ways that EMT does not. Simply put, the empirical evidence on EMT should not rule out further study of technical analysis.

Further, we have seen that the numerous studies of technical analysis do not always test rules in ways that are consistent with its theories. For example, almost all technical trading rules should factor in volume, not only price, though this is rarely done. Most technical analysis suggests that one should not be in the market at all times, though most studies ignore this. What is needed is essentially an extension of the work of Neftci (1991), in the sense that we must attempt to formalize rules that are more complex, and not simply abandon them in favor of rules that are more easily formalized. Work being done in neural nets and artificial intelligence is a step in the right direction.

To conclude, given the importance of energy futures markets, as well as the importance placed on technical analysis in the financial community, a study combining the two is certainly in order. Again, it should be noted that the present chapter is by no means an attempt to prove the validity of technical analysis. This chapter has simply attempted to show that there is enough debate within the recent empirical findings to warrant further study of the issue. Similarly, the goal the previous chapter was to show that within the theory there is a similar debate, again warranting further study of the subject. Having hopefully established this opinion convincingly, this paper will proceed in the following chapters to formally test some of the theories of technical analysis.

CHAPTER 4 – STATISTICAL METHODOLOGY

4.1 Introduction

This chapter outlines the statistical methodology to be used in testing technical analysis and market efficiency. In terms of testing for market efficiency, this chapter follows the taxonomy of Campbell et al. (1997). As such, Section 4.2 introduces tests of Random Walk 1 (RW1), while Sections 4.3 and 4.4 outline the tests of Random Walk 2 (RW2) and Random Walk 3 (RW3), respectively. Much recent work has dealt with the unit root hypothesis, and its implications for market efficiency - Section 4.5 will examine this topic. Finally, Section 4.6 summarizes the methodology to be used.

It should be noted that each of the tests that will be discussed in this chapter (indeed *any* statistical tests) embeds certain assumptions. Therefore, each test will allow us to comment on market efficiency *given* those assumptions. Of particular interest is the assumption of linearity in the data-generating process. Referring back to Table 2.1, we see that the RW3 and the martingale models *do* impose linearity, whereas RW1 and RW2 do not. As concerns the issue of technical analysis, we know, from Neftci (1991), that we must relax the linearity assumption since technical analysis is in fact based on the very premise of non-linearity. Thus, tests of RW3 do not allow us to comment on the effectiveness of technical analysis. Further, due to Campbell et al. (1997; p.32), we know that “the assumption of [RW1] is not plausible for financial prices over long time spans”. Unfortunately, we are *also* told that “testing for [RW2] is quite difficult” and that (p.41) “[placing] no restrictions on how the marginal distributions of the data can vary through time, it becomes virtually impossible to conduct statistical inference”. Therefore, while

we can carry out sophisticated tests of the random walk hypothesis, the only tests that allow us to comment on technical analysis are tests of technical analysis itself.

4.2 Tests of Random Walk 1

The various versions of the random walk hypothesis are outlined in Section 2.3.3.

By way of summary, the RW1 is characterized as follows:

$$P_t = \mu + P_{t-1} + \varepsilon_t; \quad \varepsilon_t \sim \text{IID}(0, \sigma^2)$$

In short, today's price will equal yesterday's price plus an expected change in price, the drift term μ . In addition, today's price will reflect a random shock or increment. The key to RW1 lies in the assumption of the increment being distributed identically and independently with mean zero and variance σ^2 . Essentially, "independence implies not only that increments are uncorrelated, but that any nonlinear functions of the increments are also uncorrelated" (Campbell et al. (1997), p.32). As such, the RW1 is a more restrictive case than the martingale.

As outlined in Section 4.1, it has been pointed out that RW1 is simply not plausible due to practical considerations. Nonetheless, tests of RW1 are among the earliest tests of the random walk hypothesis and *do* provide some insight into the nature of a given price series. This section outlines autocorrelation coefficients and the Ljung-Box statistic.

4.2.1 Autocorrelation Coefficients

One of the easiest ways to test the random walk hypothesis is to test for serial correlation. If the increments are shown to be correlated, this is clearly a violation of the random walk. The autocorrelation coefficient is defined as follows:

$$\rho(k) = \text{Cov}[P_t, P_{t+k}] / [(\text{Var}(P_t))^{1/2}(\text{Var}(P_{t+k}))^{1/2}]$$

Note that the autocorrelation coefficient is simply the time-series counterpart to the usual correlation coefficient between two random variables. Basically, we would expect zero autocorrelation in the logged first differences of the series of closing prices, which would imply that *returns* are uncorrelated.

Estimating the autocorrelation coefficient implies deriving the sampling theory for $\rho(k)$, which in turn implies making assumptions about the data-generating process of P_t . Under the assumptions of RW1, Fuller (1976) shows that the sample autocorrelation coefficients are asymptotically distributed as follows:

$$\rho(k) \sim N(0, 1/(T)^{1/2})$$

Note that T refers to the sample size, and that the above sample distribution gives us a means to determine standard deviations and therefore to test the significance of calculated $\rho(k)$. Again, significant autocorrelation leads to a rejection of the random walk model.

4.2.2 Ljung-Box Q-Statistic

This statistic can be used to test the hypothesis that all autocorrelations are zero. The Ljung-Box Q is in fact a finite sample correction of the earlier Box-Pierce Q -statistic. Ljung and Box (1978) suggest the use of the following statistic:

$$Q = T(T+2) \sum_{k=1}^m [\rho^2(k) / (T-k)] ; \quad Q \sim \chi^2_m$$

Note that $\rho^2(k)$ is the squared k -th autocorrelation. By summing squared autocorrelations, this statistic detects any significant deviations from zero autocorrelation, whether positive or negative. Again, zero autocorrelation in the logged first differences would imply that returns are uncorrelated. An important consideration involves the choice of number of autocorrelations m . Following Campbell et al. (1997, p.67), this study reports the Q statistics for five and ten autocorrelations. Clearly, a Q -value above the critical value implies that autocorrelation is a problem, and that the random walk does not hold.

4.3 Tests of Random Walk 2

The RW2 is seen as a more plausible model than the RW1, and can be characterized as follows:

$$P_t = \mu + P_{t-1} + \varepsilon_t ; \quad \varepsilon_t \sim \text{INID} (0, \sigma^2)$$

The key here is clearly in the way we specify the increments ε_t to be distributed. Specifically, we require the increments to be independently but not identically distributed (INID). While this is clearly a more realistic assumption, it unfortunately does not lend itself well to statistical testing. Campbell et al. (1997) suggest that while most statistical tests of RW2 are impractical if not impossible, there exist *economic* tests of RW2, including filter rules and technical analysis.

It is unfortunate that statistical tests are unavailable for the RW2. In terms of *testing* technical analysis, we are only able to test particular rules. Results based on these tests can then only tell us whether these particular rules were effective or not. Indeed, this is precisely the problem we face when we lack rigorous statistical tests. Therefore, using technical analysis we can never be sure a given series *does* truly exhibit RW2. We can only assert that the series is consistent with RW2 around the specific non-linear model that we have tested. Likewise, though many technical trading rules may prove to be consistent with RW2, we cannot rule out the possibility that other rules (perhaps ones that are more complex or better specified) will prove inconsistent with RW2.

Despite these drawbacks, we shall proceed with tests of certain trading rules. It has been mentioned in an earlier section that many of the more commonly tested trading rules are overly simplistic, leaving out key elements of technical analysis theory, such as volume and confirmation. Unfortunately, many of the more complicated trading rules have been shown to be difficult to formalize (see Neftci (1991)). Without getting into artificial intelligence and neural nets, we are somewhat limited as to the complexity of rules we can test. Therefore, this section will examine two of the more commonly used trading rules – the moving average crossover and the channel rule.

4.3.1 Moving Average Crossover

The moving average crossover is one of the simplest and most commonly used of the so-called *trend-following* technical trading rules. In its most basic form, this rule involves only the price series and a single moving average. When prices move up through the MA, a buy signal is generated. Likewise, when prices fall below the MA, a

sell signal is given. This type of system can be specified such that the trader takes only long positions, or that the trader is *in the market* always with either a long or short position.

A variation on the rule, which is said to detect trend changes more quickly, involves the use of two moving averages. Given a series of prices, we construct two moving averages of differing time lengths. The shorter period MA will be the more volatile of the two – conversely we may say that the slower (or longer period) MA will be smoother. Under this method, a buy signal is generated when the fast MA moves up through the slow MA, and vice-versa for a sell signal. Again, the trader may take long or short positions, or both.

The most important choice in this sort of system involves the choice of lengths of MA to be used. Many traders swear by short time periods, such as 4-day or 9-day, while others advocate using a 50-day with a 200-day. Often the length of MA corresponds to the length of the trading week or month (in trading days), or to seasonal effects. The optimal length of MA is sure to depend on the commodity or asset in question, as well as on the trader's preferences.

The methodology to be used in this study is derived from that presented in *Equis MetaStock*, one of the more popular technical analysis software packages. *MetaStock* enables the user to specify whether long and/or short positions (or both) are to be taken. In addition, the program will optimize the lengths of MA that should be used, in the sense that it will calculate and rank returns for all MA lengths, within parameters set by the user. While data mining could be a problem, with some experimenting, we should be

able to yield some interesting conclusions as to the effectiveness of the moving average crossover rule.

4.3.2 Channel Rule

The channel rule is a somewhat more sophisticated trading rule than the MA crossover. While there are numerous types of channel rules, they are generally price momentum indicators. The idea is that price movements, whether trending or otherwise, consist primarily of random movements or noise. We therefore construct a channel, within which price movements are considered to be insignificant. Only when prices move outside the channel is a signal generated.

This study tests the *commodity channel index* developed by Lambert (1980), as outlined in Colby & Meyers (1988). The commodity channel index (CCI) can be represented as follows:

$$CCI = (M - Mbar) / (0.15Dbar)$$

where:

$$M = 1/3 (\text{high} + \text{low} + \text{close})$$

Mbar = the n-period simple moving average of M

$$Dbar = 1/n \sum_{t=1}^n |M_t - Mbar|$$

Simply put, this indicator is designed to detect price movements which represent significant departures from average price behavior. By construction, the majority of price movements will fall within a $\pm 100\%$ channel. Movements outside the channel generate buy or sell signals.

Again using *MetaStock*, we can specify the parameters to be used in a trading rule based on the CCI. As with the MA crossover rule, while we may experiment with the parameters to find more effective trading rules, we must avoid the temptation to mine the data. For example, to find that the CCI with a 641% channel was *very* profitable between October 1 1997 and October 4 1997, would *not* be a powerful result. Nonetheless, by experimenting with the CCI under different parameters, we should be able to learn something about the effectiveness of this trading rule. This study uses a $\pm 250\%$ channel.

4.4 Tests of Random Walk 3

The RW3 is the weakest version of the random walk hypothesis, requiring only that the increments be uncorrelated. Formally, this means relaxing the assumption of independence. The RW3 can be characterized as follows:

$$P_t = \mu + P_{t-1} + \varepsilon_t, \quad \text{Cov} [\varepsilon_t, \varepsilon_{t-k}] = 0 \text{ for all } k \neq 0$$

Note that RW3 contains RW1 and RW2 as special cases. For example, simple non-linear dependence such as $\text{Cov} [\varepsilon_t^2, \varepsilon_{t-k}^2] \neq 0$ satisfies RW3 but not RW1 or RW2, due to the violation of independence.

Referring to Table 2.1, we see that the RW3 assumes linearity in the data generating process. Again, due to Neftci (1991), we know that tests of RW3 will therefore not allow us to make comments or conclusions on technical analysis, given the latter requires relaxing the assumption of linearity. Nonetheless, tests of RW3 are among the most commonly used tests of market efficiency, and are therefore of considerable interest. This section will outline variance ratios as tests of RW3.

4.4.1 Variance Ratios

One of the most common tests of the random walk hypothesis involves the variance ratio, essentially a test for homoskedasticity. The variance of P_t plus the variance of P_{t+k} should be equal to *twice* the variance of P_t . More generally, the ratio:

$$VR(q) = \text{Var}[P_t(q)] / q\text{Var}[P_t]$$

should be statistically equal to one, given q -period variance. The sampling distribution of $VR(q)$ must be derived under specific assumptions about the evolution of P_t .

Lo and MacKinlay (1988) show that the variance ratio can be estimated using the following formula:

$$VR(q) = 1 + 2 \sum_{k=1}^{q-1} (1-k/q) \rho(k)$$

Though the authors develop a more rigorous sampling distribution of $VR(q)$ under the assumptions of RW3, the above formula serves as a good approximation. Given the autocorrelation coefficients, then, we can proceed to test whether or not $VR(q)$ is significantly different from one. If it is, we reject the random walk model. As with the tests of autocorrelation coefficients and Q statistics, we test using the logged first differences of the series of closing prices, meaning that we are testing for zero autocorrelation in returns.

4.5 Unit Root Tests

All versions of the random walk and martingale hypotheses require a unit root in the level of the price series. It should be noted, however, that the presence of a unit root

is a necessary but not sufficient condition for these models to hold. Therefore, testing the unit root hypothesis is of some interest in terms of efficient markets theory.

The concept of the unit root is closely tied to that of stationarity. A stationary variable is one which is mean-reverting, meaning that shocks to it will be temporary in nature. Conversely, a non-stationary variable does *not* revert to a particular mean, and therefore shocks to this sort of variables will have permanent effects. Clearly, the random walk models will require that the variable be non-stationary, such that at any point in time the expected change in the variable will be zero, and not a reversion towards a stationary mean.

4.5.1 Augmented Dickey-Fuller (ADF) Test

In order to explain the concepts involved in the ADF test, consider the following first-order autoregressive, or AR(1) process:

$$\Delta y_t = \mu + \rho y_{t-1} + \varepsilon_t; \quad \varepsilon_t \sim \text{IID}(0, \sigma^2)$$

Note that Δy_t represents the first-difference of the series y_t (i.e. $\Delta y_t = y_t - y_{t-1}$). This AR(1) process is stationary if $-1 < \rho < 1$, whereas if $\rho = 1$ the above equation describes a random walk with drift μ and is therefore non-stationary. In other words, with $\rho = 1$, tomorrow's price is today's price, plus an expected price change (the drift μ) and a random increment ε_t – the series has infinite memory in that shocks are permanent. On the other hand, if $-1 < \rho < 1$, the series has a finite memory in that it is mean-reverting. If the absolute value of ρ exceeds one, the series is explosive.

The usual way to test the null hypothesis of a unit root is to rearrange the above equation as follows:

$$\Delta y_t = \mu + \gamma y_{t-1} + \varepsilon_t; \quad \gamma = \rho - 1$$

The null hypothesis becomes $H_0: \gamma = 0$. To reject the null hypothesis is to reject that the series has a unit root, and to reject that the process is a random walk. Clearly, the existence of a unit root is a necessary condition for the random walk to hold.

More generally, an ADF test involves regressing the first difference of a time series against the lagged series, lagged differences, and (optionally) a constant term and time trend. Incorporating all of these elements, we have the regression:

$$\Delta y_t = \alpha_1 y_{t-1} + \alpha_2 \Delta y_{t-1} + \alpha_3 \Delta y_{t-2} + \alpha_4 + \alpha_5 t$$

In this case, the unit root null hypothesis is $H_0: \alpha_1 = 1$. In general, we rearrange the regression such that $\gamma = \alpha_1 - 1$, so that our null again becomes $H_0: \gamma = 0$. The unit root null is then tested against a stationary alternative.

Specifically, in this study, we test the following ADF regressions. In order to test the hypothesis that a single unit root exists within the logarithm of the series, we test the following:

$$\Delta \log x_t = \alpha_0 + \alpha_1 t + \alpha_2 \log x_{t-1} + \sum_{i=1}^m \beta_i \Delta \log x_{t-i} + \varepsilon_t$$

In order to test the null hypothesis that a second unit root exists within the logarithm of the series, or that a unit root exists within the first differences of the logarithm of the series, we test the following:

$$\Delta \Delta \log x_t = \alpha_0 + \alpha_1 t + \alpha_2 \log x_{t-1} + \sum_{i=1}^m \beta_i \Delta \Delta \log x_{t-i} + \varepsilon_t$$

In both cases, the unit root null hypothesis is rejected if α_2 is negative and significantly different from zero. Testing this hypothesis involves comparing a t -statistic against its critical values. It is important to note, however, that standard critical values

for t do not apply, since we are dealing with a non-stationary variable in γ . Dickey and Fuller, in 1976, derived and calculated the critical t values for unit root tests. MacKinnon (1991) derived a more flexible means of calculating these values, permitting the calculation of critical values for any sample size and for any model specification. (EViews incorporates MacKinnon's critical values).

Note that the ADF regression reduces simply to the DF regression with $m=0$. Note also that for the ADF test the choice of number of autocorrelations, m , is somewhat arbitrary, though it should be large enough to ensure that the increments ε_t are white noise. Said and Dickey (1984) show that the ADF is valid if the order of the autoregression is increased with the sample size T at the rate $T^{1/3}$

4.6 Summary

This chapter has outlined the methodology to be used in testing the theories of technical analysis and market efficiency. Section 4.2 outlines autocorrelation coefficients and Q-statistics as tests of RW1. The moving average crossover and commodity channel index are presented in Section 4.3 as tests of RW2. Variance ratios, outlined in Section 4.4, provide a test of RW3. Finally, the unit root hypothesis can be tested using the Augmented Dickey-Fuller test, as outlined in Section 4.5.

Together, these tests should provide us with a great deal of information about the movement of energy prices. To the extent that the various random walk models hold, this would provide evidence in support of the EMT. In addition, those same results would

imply that trading rules should be ineffective. Taken together, these results should answer questions as to the predictability of energy prices.

An insight that should not be lost is one regarding our assumptions underlying the data generating process. Under most statistical tests, we impose linearity, which is rather restrictive and limits the sort of conclusions we may reach. When we relax linearity, such as with technical trading rules, we are really *only testing those rules*. Our conclusions can be no more powerful than that the given rule works or does not – we cannot make powerful conclusions about the random walk hypothesis using these sorts of tests. Simply put, our assumptions are of necessity quite restrictive, and this fact must be recognized when we draw our conclusions.

CHAPTER 5 – DATA & RESULTS

5.1 Introduction

This chapter proceeds with the testing of the various random walk and unit root hypotheses. Section 5.2 presents the data to be used, including information regarding its source, graphical representations of the series, as well as summary statistics of the returns generated by the series. The results from the tests of RW1 through RW3 are presented in sections 5.3 through 5.5. It should be noted that in each case the choice of which results are reported is somewhat subjective. In general, the choice is made in a way which gives results which are compatible, and therefore comparable, to those in previous studies. Section 5.6 reports the results from the various unit root tests, while Section 5.7 summarizes the results.

5.2 Data

This study uses daily one-month closing prices for six energy futures markets – light crude oil, natural gas, heating oil, unleaded gasoline, liquid propane and electricity. From the theory of technical analysis, we know that the choice of price series should not be relevant. The technical analyst claims that technical trading rules should be applicable to *any* price series, whether stocks or bonds or commodities, future or spot prices.

The decision to use closing prices of one-month futures contracts, however, seems to be logical for several reasons. One-month futures prices tend to converge on spot prices (see, for example, Fama (1984)), and can therefore be used as a proxy for spot prices. In addition, one-month futures, or *spot contracts*, tend to trade more actively than

more distant contracts. Finally, by *rolling over* contracts each month, we can obtain a continuous price series. Together, these factors make the spot contract of most interest to traders, and therefore of most interest in terms of testing technical trading rules.

Many of the energy futures have relatively limited histories. This is especially true in the cases of electricity and natural gas, due to recent deregulation in these industries. The implication of this is that each contract has a different available sample range over which we can test. While it would be preferable to have longer sample periods in some cases, we must make do with what data *do* exist.

The following table outlines the details of each contract to be used. Note that in each case the one-month contract is *rolled over*, meaning that the contract trades until the last day of the month prior to delivery, and then is replaced by the following month's contract, making for a continuous series. Charts 5.1 through 5.6 plot the logged daily closing prices over their sample range:

Table 5.1 – Details regarding the data, including sample range.

Contract	Symbol	Exchange	Start Date	End Date	Observations
Crude Light Oil	CL	NYMEX	Mar. 30, 1983	Jan.23, 1998	3687
Heating Oil	HO	NYMEX	June 02, 1980	Jan.23, 1998	4392
Liquid Propane	PN	NYMEX	Aug. 21, 1987	Jan.23, 1998	2612
Natural Gas	NG	NYMEX	April 03, 1990	Jan.23, 1998	1950
PV Electricity	EV	NYMEX	April 01, 1996	Jan.23, 1998	447
Unleaded Gas	HU	NYMEX	Dec. 03, 1984	Jan.23, 1998	3292

Chart 5.1 – Logged Daily Closing Prices for Crude Oil Spot Contracts.

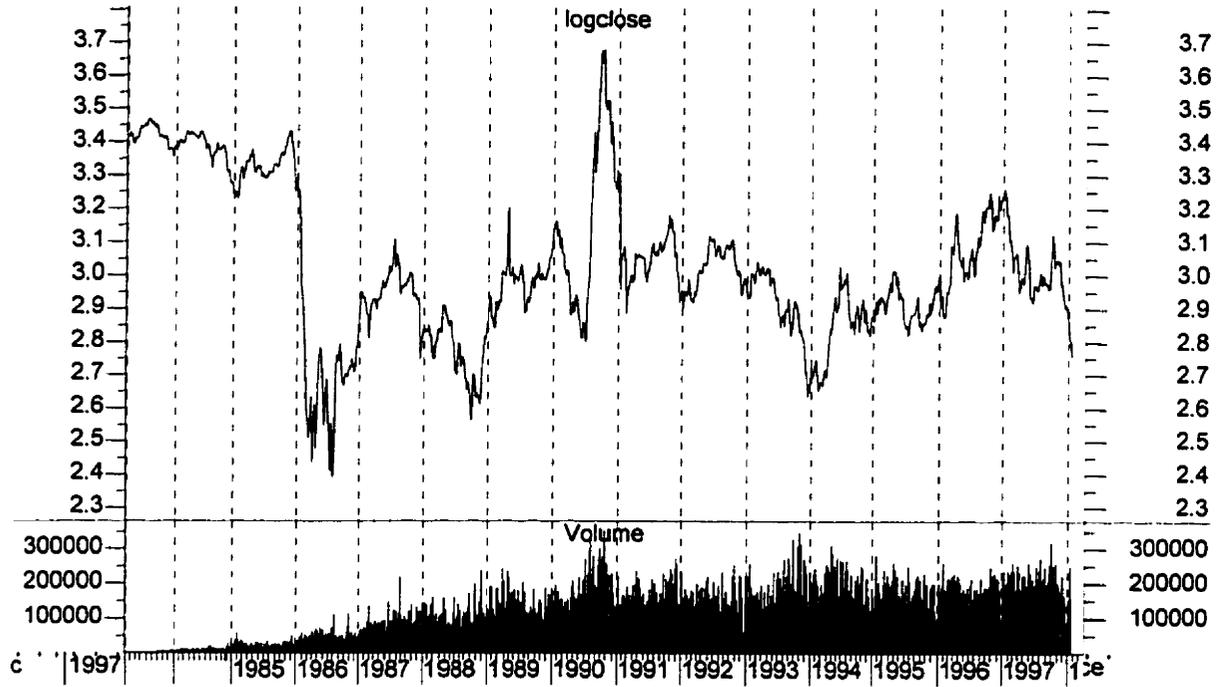


Chart 5.2 – Logged Daily Closing Prices for Heating Oil Spot Contracts.

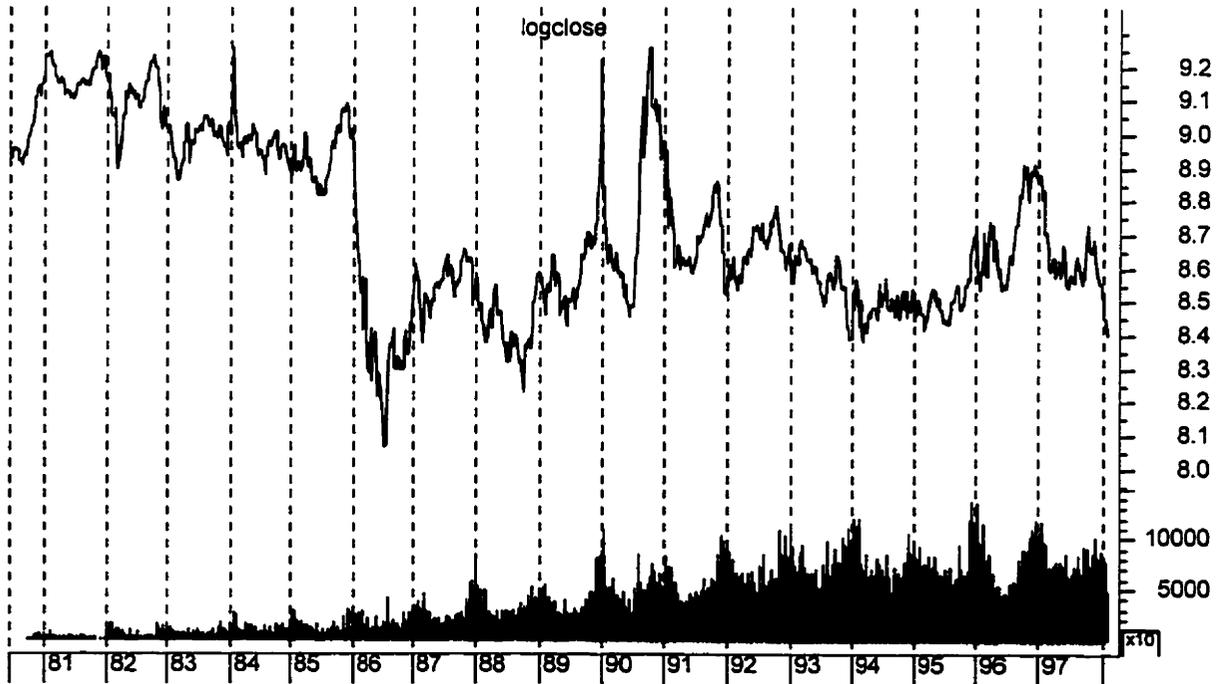


Chart 5.3 – Logged Daily Closing Prices for Liquid Propane Spot Contracts.

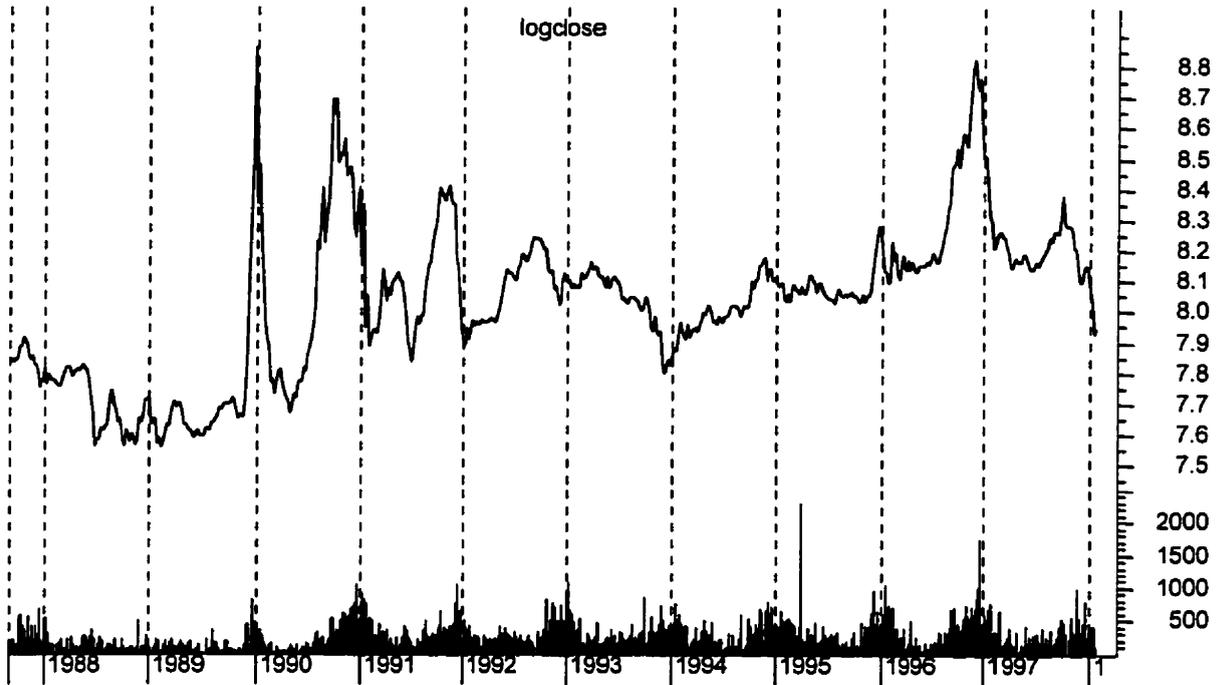


Chart 5.4 – Logged Daily Closing Prices for Natural Gas Spot Contracts.

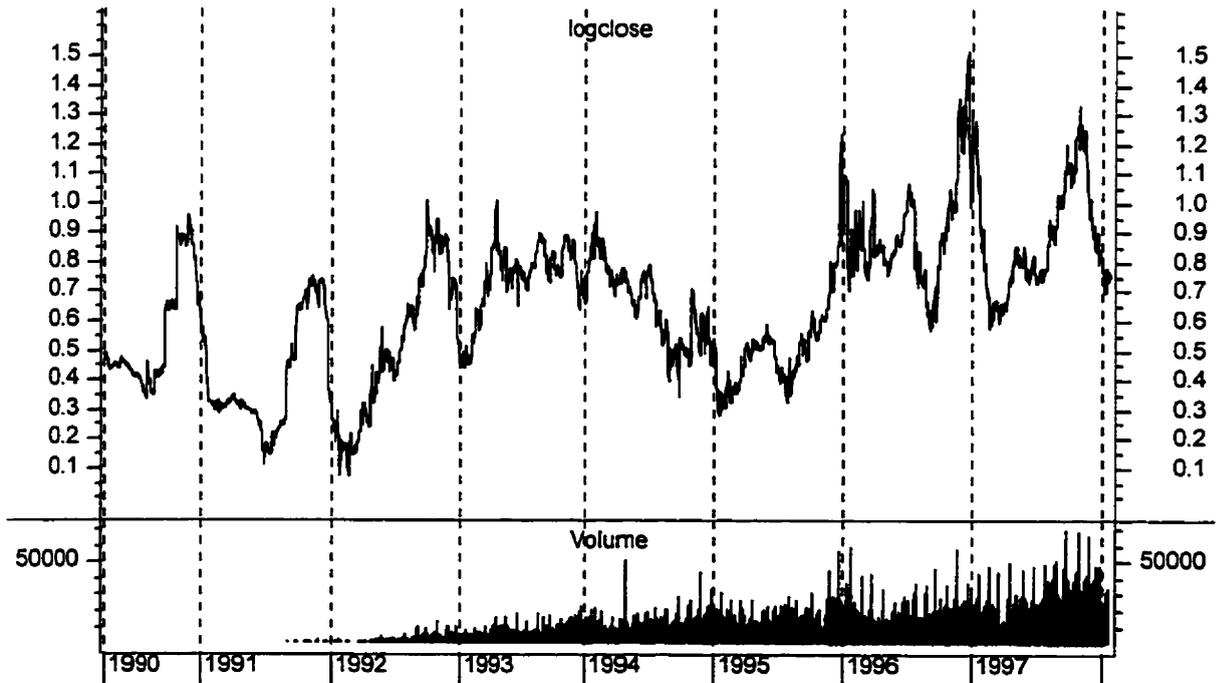


Table 5.2 presents some summary statistics for daily percentage returns in energy futures. The skewness numbers suggest that the distribution of returns is consistent with symmetry. The kurtosis numbers, however, suggest that there are too many large variations in returns for the series to be consistent with normality.

Table 5.2 – Summary statistics for daily percentage returns in energy futures.

Contract	Mean	Std. Dev.	Minimum	Maximum	Skewness	Kurtosis-3
Crude Light Oil	-1.82 E-4	0.0232	-0.4004	0.1403	-1.7653	34.593
Heating Oil	-1.26 E-4	0.0226	-0.3909	0.1399	-2.4437	37.939
Liquid Propane	3.93 E-5	0.0315	-0.7133	0.7231	-0.9016	223.77
Natural Gas	1.44 E-4	0.0350	-0.3757	0.2462	-0.5640	15.089
PV Electricity	1.41 E-3	0.0513	-0.3255	0.2905	-1.0096	13.456
Unleaded Gas	-1.30 E-4	0.0227	-0.3099	0.1234	-1.157	18.280

5.3 Tests of RW1

Table 5.3 reports the results from the tests of the RW1 hypothesis. Included are autocorrelation coefficients up to four lags, Q-statistics for autocorrelations of five and ten period lags, as well as sample size. Given these statistics, we can proceed with the testing of the RW1 hypothesis. Recall from chapter 4 that for relatively large sample sizes the sample autocorrelation coefficients are asymptotically independent and normally distributed with mean zero and standard deviation $1/\sqrt{T}$, where T is the sample size. Also, recall that the Q-statistic is distributed as χ^2 with m degrees of freedom, where m is the number of autocorrelations. Note the following $\alpha=5\%$ critical values of $\chi^2_5=11.070$, $\chi^2_{10}=18.307$ and $\alpha=2.5\%$ values of $\chi^2_5=12.833$, $\chi^2_{10}=20.483$

Table 5.3 – Autocorrelation in daily percentage returns in energy futures.

Contract	Sample Size	$1/\sqrt{T}$	$\hat{\rho}_1$	$\hat{\rho}_2$	$\hat{\rho}_3$	$\hat{\rho}_4$	Q_5	Q_{10}
Crude Oil	3687	0.0165	-0.001	-0.061	-0.082	0.044	54.015	71.238
Heating Oil	4392	0.0151	-0.010	0.014	-0.081	-0.022	38.853	64.009
Propane	2612	0.0196	-0.083	0.028	-0.042	-0.004	25.056	46.558
Natural Gas	1950	0.0226	0.015	-0.077	-0.038	0.002	14.783	18.763
PV Electricity	447	0.0473	0.014	-0.026	-0.118	0.059	8.692	10.796
Unleaded Gas	3292	0.0174	0.047	0.009	-0.081	-0.005	29.780	50.671

Looking first at the autocorrelation coefficients, we see that in the majority of cases the autocorrelations are statistically significant. For example, in the case of crude oil, we see that the second, third and fourth-order autocorrelation coefficients are significant. In the case of each series we find significant autocorrelation at some level, and we therefore reject the RW1 hypothesis.

Turning to the Q-statistics, and using the 5% level of significance, we see that only the electricity series is consistent with RW1. Relaxing the level of significance to 2.5%, the natural gas series is also consistent with RW1. From a visual inspection of Table 5.3, we see that these results seem to confirm those found using the autocorrelation coefficients. Both the electricity and natural gas series exhibited only minor deviations from zero autocorrelation, as opposed to the four other series. Note that the Q-statistic tests *jointly* that all autocorrelations are zero, and is therefore a more stringent test than the autocorrelation – this likely accounts for the slight discrepancy in the results.

Therefore, the results generally show that the series are *not* consistent with the RW1 hypothesis. Campbell et al. (1997, p.66) explain the *economic* significance of results of the type shown in Table 5.3. They point out that “the R^2 of a regression of returns on a constant and its first lag is the square of the slope coefficient, which is simply the first-order autocorrelation”. Therefore, for the case of propane, first-order autocorrelation of 8.3% means that 0.7% of variation in daily returns is predictable using the previous day’s return. Clearly, these results show a deviation from RW1 that is statistically significant, though not economically significant.

5.4 Tests of RW2- Results of Technical Trading Rules

5.4.1 Commodity Channel Index (CCI)

The results of the CCI trading rule are given in tables 5.4 through 5.9. The most striking conclusion that can be reached, from simply a visual inspection of the results, is that none of the rules performs consistently well. In fact, it would seem that most of these rules are money-losers.

Despite the generally poor results, a number of interesting observations can be made. First of all, there *are* certain combinations of channel widths which yield excess returns in certain commodities. However, a close examination of the results suggests that this may simply be the result of data-mining. For example, looking at the results for crude oil we see that a 25-275 channel system would yield a gain of over 200%, *yet* a 20-275 system would yield *losses* of almost 75%. Clearly, the results would seem to be spurious with respect to any underlying theory.

The potential for data-mining, and for coming up with rules that yield eye-popping returns, is an important issue in technical analysis. If technical trading rules are to be considered legitimate, the full range of results should be reported. The advantage of using an optimization procedure, such as the one used in this study, is that it allows us to view the profitable rules alongside the losing rules. Within this broader context, we gain a better idea of the risks associated with following a mechanical trading rule of this sort. Also to be noted is the variable returns that a single rule can yield over different series. For example, while the 20-250 system yielded 109% profits in crude oil, that same system lost over 67% on heating oil. The moral of this story – don't bet the farm on a commodity channel index trading rule!

Table 5.4 – CCI Trading Rule Results for Liquid Propane

% Gain	# Trades	Winning	Losing	Opt. 1	Opt.2
75.79	4	3	1	15	275
18.99	2	1	1	15	300
-22.23	8	6	2	25	275
-31.63	1	1	0	25	300
-36.82	16	10	6	15	225
-51.21	8	6	2	15	250
-57.06	6	3	3	20	275
-60.15	4	2	2	20	300
-65.78	16	11	5	20	250
-66.26	9	5	4	25	250
-78.70	29	15	14	25	200
-82.64	21	13	8	25	225
-93.77	29	13	16	20	200
-106.62	11	7	4	15	200
-130.92	5	3	2	20	225

Table 5.5 – CCI Trading Rule Results for Natural Gas

% Gain	# Trades	Winning	Losing	Opt. 1	Opt.2
318.63	10	8	2	15	250
171.06	19	13	6	20	225
104.23	6	4	2	20	225
22.88	8	4	4	25	250
21.09	6	3	3	20	275
-4.29	9	6	3	20	250
-9.59	6	3	3	25	275
-13.65	15	11	4	15	225
-15.33	2	1	1	15	225
-16.20	15	8	7	25	225
-23.11	2	0	2	20	300
-23.11	2	0	2	25	300
-77.40	21	12	9	20	200
-80.15	21	13	8	15	200
-90.80	17	7	10	25	200

Table 5.6 – CCI Trading Rule Results for Crude Oil

% Gain	# Trades	Winning	Losing	Opt. 1	Opt.2
202.17	13	8	5	25	275
108.96	19	12	7	20	250
94.76	5	4	1	25	300
55.62	26	18	8	20	225
30.91	27	19	8	25	225
29.34	15	9	6	25	250
18.09	39	25	14	15	225
-6.23	5	2	3	15	275
-20.84	3	2	1	15	300
-29.39	19	11	8	15	250
-46.10	40	26	14	25	200
-68.27	42	26	16	20	200
-74.96	13	8	5	20	275
-84.33	5	3	2	20	300
-85.24	52	28	24	15	200

Table 5.7 – CCI Trading Rule Results for Heating Oil

% Gain	# Trades	Winning	Losing	Opt. 1	Opt.2
517.91	30	23	7	25	225
360.42	20	15	5	25	250
122.30	54	38	16	25	200
-21.09	9	6	3	25	300
-21.18	6	4	2	15	300
-24.67	18	10	8	15	250
-29.63	14	8	6	25	275
-49.07	7	3	4	15	275
-52.23	14	6	8	20	275
-52.55	11	7	4	20	300
-63.63	58	38	20	20	200
-66.00	60	37	23	15	200
-66.22	34	21	13	15	225
-67.18	22	13	9	20	250
-79.55	34	24	10	20	225

Table 5.8 – CCI Trading Rule Results for Palo Verde Electricity Futures.

% Gain	# Trades	Winning	Losing	Opt. 1	Opt.2
297.45	2	2	0	20	275
154.78	3	2	1	20	250
45.66	0	0	0	20	300
45.66	0	0	0	20	300
3.81	0	0	0	25	300
-66.38	0	0	0	25	275
-92.60	3	0	3	25	250
-96.23	5	2	3	25	225
-98.10	5	1	4	25	200
-98.29	5	2	3	20	225
-99.84	5	0	5	20	200
-100.52	5	1	4	15	225
-100.53	5	1	4	15	200
-105.30	3	0	3	15	275
-117.94	5	3	2	15	250

Table 5.9 – CCI Trading Rule Results for Unleaded Gasoline Futures.

% Gain	# Trades	Winning	Losing	Opt. 1	Opt.2
429.47	13	11	2	20	250
389.83	26	19	7	25	225
188.17	43	31	12	25	200
86.86	14	9	5	15	275
83.09	4	3	1	15	300
61.09	31	20	11	20	225
43.49	14	8	6	15	250
12.47	45	28	17	20	200
4.14	5	2	3	25	275
-23.48	30	15	15	15	225
-30.18	2	1	1	20	300
-34.60	46	27	19	15	200
-37.58	12	7	5	25	250
-40.04	3	0	3	25	300
-49.92	6	3	3	20	275

5.4.2 Moving Average (MA) Crossover with Optimization

The results from the MA crossover tests are reported in tables 5.10 through 5.16.

At first glance, the results seem similar to those of the CCI tests – the results seem rather erratically distributed, with no consistently profitable rules. A closer inspection, however, yields some interesting observations.

Table 5.10 – MA Crossover Results for Unleaded Gasoline.

% Gain	# Trades	Winning	Losing	Opt.1	Opt.2
-11.76	584	167	417	1	10
-75.06	270	82	188	5	10
-88.85	285	57	228	1	50
-93.20	139	35	104	5	50

Table 5.11 – MA Crossover Results for Palo Verde Electricity.

% Gain	# Trades	Winning	Losing	Opt.1	Opt.2
-10.58	36	10	26	5	10
-20.27	13	6	7	5	50
-30.50	27	7	20	1	50
-52.67	74	16	58	1	10

Table 5.12 – MA Crossover Results for Natural Gas.

% Gain	# Trades	Winning	Losing	Opt.1	Opt.2
116.14	122	30	92	1	50

115.47	60	22	38	5	50
97.24	144	42	102	5	10
-28.06	346	94	252	1	10

Table 5.13 – MA Crossover Results for Liquid Propane.

% Gain	# Trades	Winning	Losing	Opt.1	Opt.2
4077.12	371	127	244	1	10
400.56	195	58	137	5	10
71.02	187	39	148	1	50
50.83	109	29	80	5	50

Table 5.14 – MA Crossover Results for Heating Oil.

% Gain	# Trades	Winning	Losing	Opt.1	Opt.2
1.71	320	79	241	1	50
-70.39	166	45	121	5	50
-72.85	778	229	549	1	10
-91.62	368	106	262	5	10

Table 5.15 – MA Crossover Results for Crude Oil.

% Gain	# Trades	Winning	Losing	Opt.1	Opt.2
-74.80	687	188	499	1	10
-75.49	145	38	107	5	50
-79.35	305	60	245	1	50
-95.35	339	90	249	5	10

First, it should be noted that the four rules tested seem to be consistently profitable *or* consistently losing - the key being *consistently*. Returns are positive for propane and natural gas, but are persistently negative for the other four series. In fact, returns are *so* poor in some cases that it suggests that the rule may be generating some sort of predictability. Indeed, to the extent that returns *are* persistently negative, beyond that caused by commissions, then perhaps there are other rules which may capture that predictability.

While it is not the intent of this study to mine the data so as to find rules which *do* seem profitable, it is important to note that there are no theoretical reasons why this could not be done. Not only is the evidence regarding the random walk lacking, but more importantly, tests of the random walk impose linearity, and therefore generate results which cannot be used in terms of commenting on technical analysis. Only tests of technical analysis can be used to comment on the subject. Therefore, the failure of any given rule is a failure of that particular rule, and is in no way any more general a statement than that. The results of the tests RW2 are of necessity inconclusive.

5.5 Tests of RW3

Table 5.17 reports the results of the variance ratio tests. Included are the variance ratios, the sample sizes and the standardized normal test statistics. Recall from Section 4.4.1 that the variance ratio can be estimated using the following formula:

$$VR(q) = 1 + 2^{q-1} \sum_{k=1}^{q-1} (1-k/q) \rho(k)$$

Note that VR should equal one under the random walk model.

The following standardized normal test statistic can be used to test the hypothesis that $VR=1$:

$$\psi(q) = \sqrt{nq} [VR(q)-1] [(2(2Q-1)(q-1)) / 2q]^{-1/2}$$

where nq is simply equal to the sample size, following Campbell et al.(1997, p.69). At the 5% level of significance, the critical p -value is 1.96.

Table 5.16 – Variance Ratios for Percentage Daily Returns in Energy Futures

Contract	Sample Size	VR(2) ($\psi(2)$)	VR(4) $\psi(4)$	VR(8) $\psi(8)$	VR(16) $\psi(16)$
Crude	3687	0.999 (-0.061)	0.897 (-3.343)	0.807 (-3.962)	0.757 (-3.352)
Heating	4392	0.990 (-0.663)	0.959 (-1.452)	0.814 (-4.167)	0.764 (-3.553)
Propane	2612	0.917 (-4.242)	0.883 (-3.196)	0.804 (-3.386)	0.839 (-1.869)
Natural	1950	1.015 (0.662)	0.923 (-1.817)	0.839 (-2.403)	0.815 (-1.856)
Electricity	447	1.014 (0.296)	0.936 (-0.723)	0.867 (-0.951)	0.893 (-0.514)
Unleaded	3292	1.047 (2.697)	1.039 (1.196)	0.966 (-0.659)	0.944 (-0.730)

The results of the variance ratio tests are interesting for a number of reasons.

Given a 5% level of significance, only the electricity series exhibits variance ratios that are unambiguously consistent with the random walk model. In each of the other series,

we see test statistics exceed their critical values for certain values of q , implying that autocorrelation is indeed a problem.

Another interesting point to note is that in every case, the tendency towards negative autocorrelation is reversed as q increases. Therefore, negative autocorrelation tends to be a factor mainly over the short run, and diminishes over time.

5.6 Unit Root Tests

Tables 5.4 and 5.5 report the results of Dickey-Fuller (DF) and Augmented Dickey-Fuller (ADF) tests. Recall from chapter 4 that these are tests of the null hypothesis that a unit root exists within the given series. Recall also that the existence of a unit root is required in order to accept non-stationarity and the random-walk hypothesis.

In order to test the null hypothesis that a single unit root exists in the logarithms of the energy futures series, we use the following ADF regression:

$$\Delta \log x_t = \alpha_0 + \alpha_1 t + \alpha_2 \log x_{t-1} + \sum_{i=1}^m \beta_i \Delta \log x_{t-i} + \varepsilon_t$$

where x_t is the particular series being tested. The null hypothesis of a single unit root in the series is rejected if α_2 is negative and significantly different from zero. Note that the ADF regression reduces simply to the DF regression with $m=0$. Note also that for the ADF test the choice of number of autocorrelations, m , is somewhat arbitrary, though it should be large enough to ensure that the increments ε_t are white noise. Said and Dickey (1984) show that the ADF is valid if the order of the autoregression is increased with the sample size T at the rate $T^{1/3}$ – these values are shown in tables 5.4 and 5.5.

The results in table 5.4 present a somewhat mixed picture. In the “without trend” version of the test, only in the electricity series do we *not* reject the unit root null with both the DF and ADF tests. In other words, the electricity series gives us the strongest evidence in favor of the random walk. However, looking at the ADF statistics only, again for the “without trend” test, and relaxing the level of significance somewhat, we would only reject the unit root null in the cases of liquid propane and unleaded gas. In other words, only in the cases of liquid propane and unleaded gas do we have reasonable evidence *against* the random walk. Allowing for trend, the picture is not much clearer. Electricity and crude oil are consistent with the random walk, whereas propane, unleaded and natural gas are not. Heating oil presents mixed results.

Table 5.17 – Tests for a unit root in the logarithms of the series.

Contract	Autocorrelations for ADF ($\sim T^{1/3}$)	Without Trend		With Trend	
		DF	ADF	DF	ADF
Crude Light Oil	15	-3.231	-2.877	-3.358	-2.976
Heating Oil	16	-2.869	-2.389	-3.649	-3.076
Liquid Propane	14	-3.365	-3.682	-4.072	-4.439
Natural Gas	12	-3.157	-2.865	-3.768	-3.422
PV Electricity	8	-2.328	-2.249	-2.231	-2.085
Unleaded Gas	15	-3.656	-3.706	-3.672	-3.714

Note: The 95% critical values for the ADF and DF test statistics (approximately equal for all lags) are – 2.863 for the “without trend”, and –3.414 for the “with trend” version of the test.

Overall, we may conclude that in log-levels, electricity and crude oil are the most consistent with a random walk. Propane and unleaded gas are the least consistent with the random walk. Natural gas and heating oil present mixed evidence. In order to test the null hypothesis that a second unit root exists – that is, a unit root in the first differences of the logged series - we use the following ADF regression:

$$\Delta \log x_t = \alpha_0 + \alpha_1 t + \alpha_2 \Delta \log x_{t-1} + \sum_{i=1}^m \beta_i \Delta \log x_{t-i} + \varepsilon_t$$

where x_t is the particular series being tested. The null hypothesis of a second unit root is rejected if α_2 is negative and significantly different from zero.

Table 5.18 – Tests for a unit root in the first differences of the logarithms of the series.

Contract	Autocorrelations for ADF ($\sim T^{1/2}$)	Without Trend		With trend	
		DF	ADF	DF	ADF
Crude Light Oil	15	-60.761	-14.977	-60.753	-14.976
Heating Oil	16	-66.914	-32.675	-66.907	-32.672
Liquid Propane	14	-55.450	-11.818	-55.491	-11.818
Natural Gas	12	-43.469	-12.274	-43.458	-12.272
Electricity	8	-20.769	-7.293	-20.776	-7.354
Unleaded Gas	15	-54.871	-13.772	-54.862	-13.770

Note: The 95% critical values for the ADF and DF test statistics (approximately equal for all lags) are -2.863 for the “without trend”, and -3.414 for the “with trend” version of the test.

Looking at table 5.5, we see that the null hypothesis of a second unit root is rejected for all series. Therefore, we may conclude that none of the series are $I(2)$. Given the evidence from table 5.4, we can say that electricity and crude oil are likely $I(1)$, meaning they have a stochastic trend. As for the other four series, the evidence is mixed, but we may speculate that they exhibit some sort of mean-reverting behavior.

5.7 Summary

This chapter has outlined the data used, as well as the results of the various tests of the random walk hypotheses. Section 5.2 outlined the details regarding the data to be used in the study. It was seen that the returns exhibited skewness numbers that were consistent with symmetry, but kurtosis numbers that suggested excess variability. Sections 5.3 through 5.5 reported the results of the tests of RW1 through RW3, while Section 5.6 reported the results of several unit root tests.

In general, the results show that returns do not follow a random walk. Only the electricity series repeatedly yielded results consistent with the random walk hypothesis. In some ways this is not surprising, given that this is the most recent series, in the sense that it only began trading in 1996. Several studies (see, for example, Campbell et al. (1997, p.69) report that predictability in returns seems to be declining over time as markets become more sophisticated and more efficient. Therefore, deviations from the random walk in the other series could be a result of predictability over the earlier portion of the sample period. An interesting extension might explore this topic.

It should be noted that statistical deviations from the random walk do not necessarily constitute *economic* deviations from an efficient market. Thus, even though

measures of autocorrelation were found to be significant, it is not clear that they represent an exploitable inefficiency in the market.

An important result is that, for the most part, our results do little in terms of resolving the debate over technical analysis. Essentially, the tests of the random walk impose linearity on the data generating process, which in itself is inconsistent with the premises of technical analysis. Therefore, our results do not allow us to resolve the issue.

CHAPTER 6 – CONCLUSIONS

This chapter provides an overview of the entire thesis, summarizing the results and highlighting the major conclusions. Referring back to the introduction, we should reflect upon the objectives of the study as they were then laid out. A primary objective was to be able to *comment* on the paradox that financial economists provide evidence showing markets to be efficient, while investors and traders are willing to invest considerable resources in attempting to outperform the market. Section 6.1 discusses this issue. A second objective was to attempt to reconcile the two competing schools of thought in order to seek out any common ground that may exist between the two. A final objective was to be able to draw on these results in order to suggest fruitful areas for future study and research. Section 6.2 deals with the former, Section 6.3 the latter.

6.1 Commenting on the Paradox

This study has shown that the paradox in question is actually much less of a paradox than might have been thought. In Chapter 2, the background theories of both technical analysis and the efficient markets theory were outlined. Through this process, it was seen that much of the debate between the two schools of thought can in fact be attributed to the different starting assumptions of each.

The approach of the financial economist, as with most sciences, is one that relies heavily on statistical inference. Unfortunately, this approach often requires that we make rather restrictive simplifying assumptions. Of primary importance, in this case, is the assumption of linearity in the data generating process. While linearity may be quite

restrictive, and perhaps unrealistic, it enables us to conduct statistical inference. The less restrictive the assumptions that we make, the less we are able to say about our data.

The approach of the technical analyst is much less reliant on mathematics and statistical inference. In fact, it often relies primarily on visual inspection of charts. While this means that the approach does not require that we make restrictive assumptions, it imposes a more serious constraint. Without statistical inference, we are severely restricted as to what we can actually prove or predict with certainty. Without testable predictions, the entire body of theories cannot be *disproven*. We are left in a sort of no-man's land, unable to reach any firm conclusions on the matter.

The paradox, therefore, can be resolved more through a study of the theory than through a study of the results. In essence, we find that we are comparing apples and oranges. Simply put, the results reached by the two schools of thought cannot be compared when the starting assumptions differ in important ways. The results are in a sense superfluous to the debate, though they are certainly helpful in terms of gaining a better understanding of how the different approaches compare.

In terms of the empirical results, the evidence did not provide resounding support for the various random walk hypotheses. In general, the tests of RW1 and RW3 showed that the series' in question did *not* follow a random walk, with returns in the electricity series being the exception. The tests of RW2 provide interesting results, but do not allow us to make definitive statements regarding the random walk model.

An important caveat is that, while the results *do* point towards statistical deviations from the random walk, it is not clear that these represent economic deviations. That is, while certain tests pointed out excessive autocorrelation, even modest transaction

costs would have been enough to erase profitability in most cases. Results of this sort do *not* throw into doubt the fundamental concepts of the efficient markets literature.

6.2 Reconciling the Theories

The financial economics / academic approach and the investment industry / technical analysis approach differ in several important ways. In the previous section, the assumption of linearity in the data generating process was discussed. Another important set of assumptions involves heterogeneous agents and rational expectations. In order to pursue rigorous statistical inference, the financial economist *must* make some simplifying assumptions in this regard. If we allow every agent to act independently, and to interpret information differently, we are left with no possibility for statistical inference. The technical analyst, on the other hand, makes few simplifying assumptions, but is unable to *prove* anything. His results are of the hit-and-miss variety, and offer untestable implications.

It is worth contemplating, however, what happens when we relax some of the assumptions. Relaxing the assumption of linearity in the data generating process is crucial. It implies discarding much of the evidence provided in the efficient markets literature, but is a more realistic assumption. With recent advances in areas such as non-linear dynamics, as well as greatly increased computing capabilities, it is foreseeable that much statistical inference will become possible *without* the assumption of linearity.

Another assumption worth relaxing would be that of homogeneity across economic agents. Again, this introduces considerable complexity into the process of statistical inference. While it would be unreasonable to introduce complete heterogeneity

into a model, allowing for two or more investor types could considerably alter the results. It is certainly a more realistic assumption that some investors will face different information sets than others, and that some investors will interpret information differently than others. This in turn involves modifying our view of rational expectations, and allowing some investors to act irrationally.

One of the dangers that arises when we proceed with unrealistic assumptions is that we miss out on important dynamics, which in fact define the system. Common sense and experience show that many investors act irrationally at times, and that some seem to consistently underperform the market. Having a model which fails to account for this fact is simply unacceptable. Examining *ex-ante* returns, assuming all agents acted rationally, and *then* concluding that markets are efficient, is faulty logic. We are left with a model that is unable to show whether we have persistent winners and losers, and one which is unable to provide realistic dynamics.

It is the view of this author that *without* relaxing some of these assumptions, the results of financial economics are of limited value. There is no doubt that the body of evidence currently available greatly enhances our understanding of financial markets. However, to be of continued relevance, the field of financial economics must evolve in ways that involve less restrictive assumptions.

Finally, in terms of reconciling the two theories, the following should be noted. The technical analyst claims that “market action discounts everything”. This is a claim that *markets* are informationally efficient, in the sense that the market as a whole quickly absorbs the information and reaches an equilibrium. It is *not* a claim that investors act rationally. It may be the case that certain investors continually misinterpret information,

while others profit from these mistakes. Once we relax the assumption of homogeneity across agents, it is not clear that the two theories have irreconcilable differences.

6.3 Directions for Future Research

The previous section showed that the major differences between the two theories lies within their starting assumptions. While these differences certainly represent a huge gulf between the theories, it has been pointed out that there *are* certain similarities. Given these similarities, it may be the case that some sharing of ideas between the two schools of thought is possible. Specifically, perhaps some of the poorly defined rules-of-thumb of technical analysis can be tested formally within the framework of EMT.

EMT assumes homogeneity while technical analysis assumes that the heterogeneity of market participants *defines* the buying and selling pressures. What happens when we relax the assumption of homogeneity? If we assume even two investor types, the dynamics change entirely. Given heterogeneity, many of the stories of technical analysis become more plausible.

When we introduce different investor types, and allow them different information sets and different behavior, our model becomes infinitely more complex. However, given recent advances in theory and computing capabilities, testing these sorts of models should become a possibility. In this way, more realistic assumptions can be admitted to our models, and more plausible dynamics can be described.

The concepts of technical analysis in many cases describe quite accurately the behavior we see around us. Investors often set arbitrary price objectives for the assets they own, and often *herd* around these price targets due to the influence of analysts'

recommendations. Support and resistance levels become plausible under a different set of starting assumptions, and these are the sorts of theories that should be tested more thoroughly. Another concept worthy of further study would be that of price patterns. Given different investor types, it is plausible that certain investors act in predictable ways, and that the interplay between the investor types could generate at least *vague* patterns. The idea here is not that the theories of technical analysis would allow for excess profits, but simply that a more thorough testing of those theories could enhance our understanding of market dynamics.

Finally, perhaps the greatest challenge is to find ways to let the EMT evolve. The current situation is that the EMT provides remarkable evidence that markets *are* efficient, despite its unrealistic assumptions. The key is to relax those assumptions, and still find ways to explain the results we have. Given this challenge, there is little doubt that the EMT will remain a topic of heated debate for many years, and that as it evolves it will continue to enhance our understanding of financial markets.

REFERENCES

- Blume, Lawrence, David Easley and Maureen O'Hara (1994). "Market Statistics and Technical Analysis: the Role of Volume". *The Journal of Finance*, XLIX (1) 153-181.
- Bopp, Anthony E. and Scott Sitzler (1988). "On the 'efficiency' of futures markets". *Energy Economics*, July 1988, pp. 199-205.
- Brock, William, Josef Lakonishok and Blake LeBaron (1992). "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns". *The Journal of Finance*, XLVII (5) 1731-1764.
- Campbell, John Y., Andrew W. Lo and A. Craig MacKinlay (1997). *The Econometrics of Financial Markets*. Princeton University Press: Princeton, New Jersey.
- Edwards, Robert D. and John Magee (1966). *Technical Analysis of Stock Trends*. Springfield, Massachusetts.
- Fama, Eugene F. (1970). "Efficient Capital Markets". *The Journal of Finance*, 25, 383-417.

Fama, Eugene F. (1991). "Efficient Capital Markets: II". *The Journal of Finance*, XLVI (5), 1575-1617.

Gjolberg, Ole (1985). "Is the spot market for oil products efficient?" *Energy Economics*, October 1985, pp. 231-236.

Herbert, John H. (1995). "Trading volume, maturity and natural gas futures price volatility". *Energy Economics*, 17 (4) 293-299.

Malkiel, Burton G. (1973). *A Random Walk Down Wall Street*. W. W. Norton & Company: New York, N.Y.

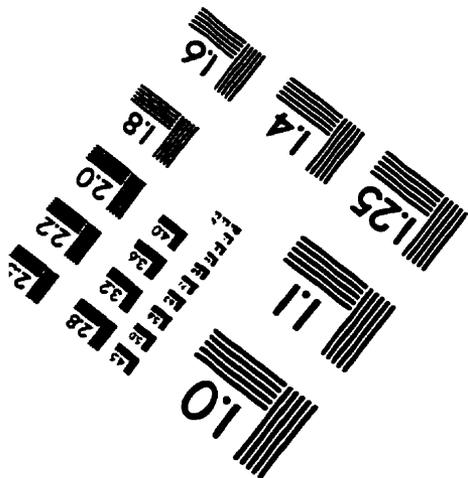
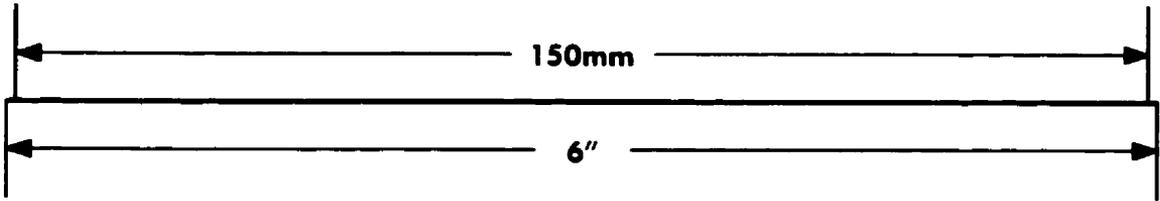
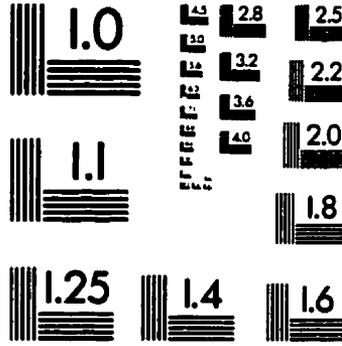
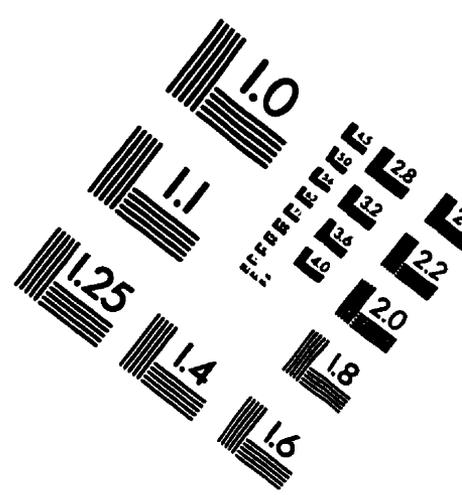
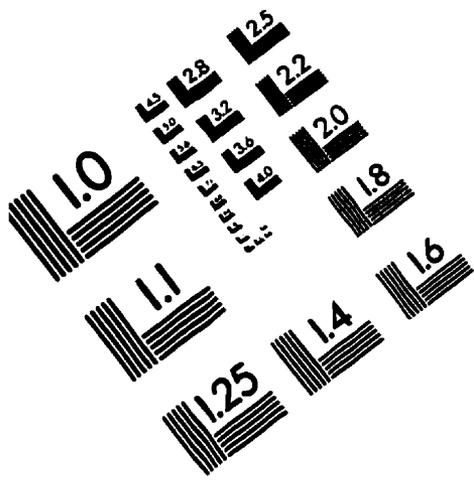
Murphy, John J. (1986). *Technical Analysis of the Futures Markets*. New York Institute of Finance: New York, N.Y.

Nainar S. M. Khalid (1993). "Market information and price volatility in petroleum derivatives spot and futures markets". *Energy Economics*, January 1993, pp. 17-24.

Neely, Christopher J. (1997). "Technical Analysis in the Foreign Exchange Market: A Layman's Guide". *Federal Reserve Bank of St. Louis Review*, September / October 1997, pp. 23-38.

- Neftci, S.N. (1991). "Naïve Trading Rules in Financial Markets and Wiener-Kolmogorov Prediction Theory: A Study of 'Technical Analysis' ". *Journal of Business*, 64, 549-571.
- Panas, Epaminondas. "A weak form evaluation of the efficiency of the Rotterdam and Italian oil spot markets". *Energy Economics*, January 1991, pp. 26-32.
- Serletis, Apostolos (1992). "Maturity effects in energy futures". *Energy Economics*, April 1992, pp. 150-157
- Serletis, Apostolos (1994). "A cointegration analysis of petroleum futures markets". *Energy Economics*, 16 (2) 93-97.
- Shiller, Robert J. (1989). *Market Volatility*. MIT Press: Cambridge, Massachusetts.
- Taylor, Mark P. (1988). "An Empirical Study of Practitioners' Predictions". *Economica*, May 1988, pp. 185-202.
- Taylor, Mark P. and Helen Allen (1992). "The use of technical analysis in the foreign exchange market". *Journal of International Money and Finance*, 11, 304-314.
- Taylor, Stephen J. (1994). "Trading Futures Using a Channel Rule". *The Journal of Futures Markets*, 14 (2) 215-235.

RESOLUTION EVALUATION TEST TARGET (QA-3)



APPLIED IMAGE . Inc
1653 East Main Street
Rochester, NY 14609 USA
Phone: 716/482-0300
Fax: 716/288-5989

© 1993, Applied Image, Inc., All Rights Reserved

