

12-7-2015

Technical Analysis under Knightian Uncertainty

Andre Mouton

CUNY Bernard M Baruch College

Follow this and additional works at: https://academicworks.cuny.edu/bb_etds

 Part of the [Economics Commons](#), and the [Finance and Financial Management Commons](#)

Recommended Citation

Mouton, Andre, "Technical Analysis under Knightian Uncertainty" (2015). *CUNY Academic Works*.
https://academicworks.cuny.edu/bb_etds/69

This Thesis is brought to you for free and open access by the Baruch College at CUNY Academic Works. It has been accepted for inclusion in Student Theses by an authorized administrator of CUNY Academic Works. For more information, please contact AcademicWorks@cuny.edu.

Technical Analysis Under Knightian Uncertainty

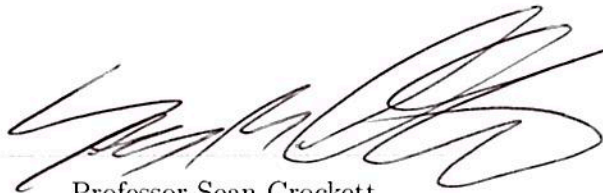
Andre Mouton

Submitted to the Committee on Undergraduate Honors at Baruch College of the City of New York on December 7, 2015 in partial fulfillment of the requirements for the degree of Bachelor of Arts in Economics with Honors.



Professor Sebastiano Manzan

Faculty Sponsor



Professor Sean Crockett

Departmental Honors Committee



Professor Guillaume Haeringer

Departmental Honors Committee

Technical Analysis Under Knightian Uncertainty

Andre Mouton*, CUNY Baruch College

December 7, 2015

Abstract

Technical analysis, or the forecasting of asset price movements using past prices, is commonly practiced in financial markets but poorly explained by mainstream economic theory. I show that a technical rule can have predictive power when an asset's payoffs are subject to Knightian uncertainty, defined as variation that cannot be described probabilistically (Knight, 1921). I present an asset-pricing model in which asset payoffs undergo periodic shifts in trend, and agents form expectations about these payoffs using a constant gain least squares (CGLS) rule. I investigate whether a second CGLS rule, operating on price, can provide a more accurate forecast of payoffs during the periods following a trend shift. I estimate the model using corporate earnings data from the S&P 500, and present simulation results that show support for the usefulness of technical analysis. Because technical analysis may influence the behavior of asset markets, this finding has potential implications for investment, risk management and financial policy.

*I would like to thank Dr. Sebastiano Manzan for his advice (and his patience), and Drs. Sean Crockett and Guillaume Haeringer for their assistance with this project.

Contents

1	Introduction	3
2	Literature Review	5
2.1	Technical Analysis	5
2.2	Knightian Uncertainty	7
3	Model	9
3.1	Stochastic Process	10
3.2	Fundamental Rule	11
3.3	Technical Rule	12
4	Results	13
4.1	S&P 500 Monthly Earnings	14
4.2	Corporate Quarterly Earnings	17
4.3	Alternative Discount Rates	19
5	Conclusion	20

1 Introduction

A commonly observed behavior in asset markets is the use of technical analysis, defined by Park and Irwin (2007) as “a forecasting method of price movements using past prices, volume, and open interest.” Surveys find that the vast majority of financial professionals report using technical analysis (Taylor and Allen, 1992; Oberlechner, 2001; Menkhoff, 2010), despite hostility from mainstream economic and financial theory. The two dominant paradigms in economics and finance, rational expectations theory and the Efficient Markets hypothesis, assert that asset prices should reflect all available information - in which case technical analysis will not be predictive. Burton Malkiel famously wrote in *A Random Walk Down Wall Street* (1973) that “technical strategies are usually amusing, often comforting, but of no real value,” and forty years later, irrationality remains the most common explanation put forward in the academic literature for the use of technical analysis (Menkhoff and Taylor, 2007). Nevertheless, empirical studies over the last three decades have generally shown technical rules to be profitable (Park and Irwin, 2007). Additionally, a number of economic models have been put forward in which the presence of technical traders is used to explain other important phenomena, such as bubbles and busts in equity markets. Technical trading may affect investment returns, economic outcomes, and policymaking decisions; and so it would be useful if we could say why, and under what conditions, this behavior is likely to occur. In this thesis, I hypothesize that technical analysis is a response to informational constraints. I construct an asset-pricing model in which traders lack probabilistic knowledge about an asset’s future payoffs - a condition often referred to as Knightian uncertainty - and I show through model simulations that a simple technical rule can be predictive in this setting.

The concept of non-probabilistic uncertainty was first introduced in an economic context by Frank Knight in *Risk, Uncertainty, and Profit* (1921). Knight distinguished between variation in economic outcomes that can be described probabilistically, or risk, and the unforeseeable and ‘uncertain’ changes in technology and taste that take place in an evolving economy. He suggested that entrepreneurs earn economic profit as a reward for contending with the latter. Similarly, John Maynard Keynes attributed cyclical investment activity to the “outstanding fact” of “the extreme precariousness of the basis of knowledge on which our estimates of prospective yield have to be made,” and the role that confidence plays in determining how individuals respond to this limitation (Keynes, 1936 12.iii). Uncertainty implies that asset prices will depend not only upon information about an asset’s payoffs, but also upon how that information is interpreted, and the confidence that market participants

place in their interpretation. In this context, technical analysis may, as suggested by its practitioners, reveal something about “the changing attitudes of investors toward a variety of economic, monetary, political, and psychological forces” (Pring, 2002).

To test this hypothesis, I construct an asset pricing model in which agents do not know the probability distribution of future dividends, but instead learn about dividends using statistical or ‘adaptive’ expectations. This approach is suggested by Evans and Honkapohja (2001) for models with incomplete information, but is nevertheless a departure from rational expectations theory, which requires that agents have full information about the model they inhabit. Economic change is introduced to this model in the form of a stochastic process of dividends that alternates periodically between states of expansion and contraction. I estimate this model using S&P 500 earnings data provided by Shiller (2015), as well as the individual earnings data of four S&P 500 components: Microsoft, GE, Exxon and Johnson & Johnson. I run simulations to test whether technical analysis - modeled as adaptive expectations about price - offers additional predictive power in periods following a state shift. The results support my hypothesis. Technical analysis is predictive for a wide range of parameter values, and performs better as the amount of Knightian uncertainty in the model increases.

The main contribution of this thesis is that it provides a generalizable explanation for the use of technical analysis. Although a handful of researchers have studied the link between technical trading, learning and uncertainty (Joshi et al, 2000; Zhu and Zhou, 2009), I present a broader framework that involves fewer behavioral assumptions. In the model I construct, all agents employ the same learning mechanism; they only differ in whether they choose to learn about dividends or about price. This consistency suggests that a different learning mechanism would produce similar results. Moreover, since there already exist a number of studies on the measurement of macroeconomic uncertainty, this thesis points the way towards empirical analyses that might further establish (or reject) the proposed relationship between technical trading and uncertainty. The theoretical model and results I discuss here should be seen merely as an initial test, and as a proof-of-concept for additional research that may benefit investors, financial institutions and policymakers.

The structure of this thesis is as follows. In Section 2 I discuss the literature on technical analysis and uncertainty. I develop a model in section 3, with estimation and simulation results given in section 4. In section 5, I discuss results and present conclusions.

2 Literature Review

2.1 Technical Analysis

Technical analysis is widely practiced in financial markets. An early survey by Smidt (1965) found that more than half of amateur traders in US commodity markets rely moderately to heavily on analysis of price charts. Brorsen and Irwin (1987) documented similar prevalence amongst advisory groups for large public futures funds, with more than half reporting heavy use of technical trading systems. A landmark study by Taylor and Allen (1992) found that 90% of London-based foreign exchange dealers employ some form of technical analysis - a number subsequently confirmed by Oberlechner (2001) - and Menkhoff (2010) found likewise that 87% of international fund managers rely on technical strategies to at least some extent. With respect to foreign exchange markets, Menkhoff and Taylor (2007) list three stylized facts: that technical analysis is widely used by professional traders, that it is typically used in conjunction with fundamental analysis (i.e. forecasts based upon an asset's expected pay-offs), and that its use increases at short time horizons.

Evidence indicates that technical strategies may be profitable. While early studies such as Fama and Blume (1966) were inconclusive, research since the late 1980's has documented profitability both in equity markets (Brock et al, 1991; Taylor, 2000) and in foreign exchange markets (Sweeney, 1986; Neely et al, 1997). A meta-analysis by Park and Irwin (2007) found that 58 of 92 studies conducted in the last three decades have shown technical rules to be profitable and/or predictive. Indeed, the annual revenues of U.S.-based high frequency trading - much of it technical in nature, or based upon non-price sentiment indicators (Manzoor, 2013) - were estimated at \$5 billion by Brogaard et al (2014) for the years 2008-2009. With respect to profitability, Menkhoff and Taylor (2007) offer an additional set of stylized facts. Technical analysis in currency markets may be profitable even after transaction and interest rate costs are taken into account, profitability tends to increase with exchange rate volatility, and the performance of any given technical rule is unstable over time. The relationship between volatility and technical profits is suggestive, insofar as volatility has also been found to correlate with political or economic uncertainty in equity markets (see for instance Bittingmayer, 1998).

The prevalence and profitability of technical analysis is difficult to explain from a theoretical standpoint. Rational expectations theory and the Efficient Markets hypothesis both

assert that prices incorporate all available information, and represent the best forecast of an asset's future performance. If this is true, then technical analysis should not provide any additional predictive power. Consequently, "the charge of not-fully-rational behaviour on the part of those applying technical analysis is probably the most common position in explaining its use" (Menkhoff and Taylor, 2007). Three other common explanations include central bank intervention, slow diffusion of market news, and the presence of other technical traders (ibid). However, technical analysis predates modern central banking in the U.S. (Dow theory originated in the 19th century), its use has been robust to changes in information technology, and a zero-sum game would imply that up to 90% of professional traders are risk seeking. A larger problem is that these explanations are non-unique and non-general. They describe idiosyncratic circumstances under which prices will convey information, whereas technical analysis is a robust phenomenon that would seem to require a more general mechanism.

Although technical analysis yet lacks a general explanation, it has proven useful in explaining the poor empirical performance of standard asset-pricing theories. For instance, Shiller (1981) demonstrates that real-world stock prices exhibit "excessive volatility," and Poterba and Summers (1988) present evidence that stock prices are mean reverting, and not a random walk as supposed by theory. De Long et al (1990) find that these anomalies can be caused by "noise traders" who "get their pseudosignals from technical analysts, stockbrokers, or economic consultants and irrationally believe that these signals carry information." Similarly, Lux and Marchesi (1998) show that a market with technical traders can produce temporally clustered volatility, matching the observed tendency for stock market volatility to show up in 'waves.'

By explaining the use of technical analysis, we might better understand its effect upon asset prices, better predict when this effect will be significant, and better account for any implications that technical analysis presents for asset-pricing theory and economic policy. There are several papers in the academic literature that analyze technical analysis in a similar framework to the one presented in this thesis. Joshi et al (2000) simulate a heterogeneous agent model with learning, and find that technical analysis is a dominant strategy. Zhu and Zhou (2009) analyze the performance of moving average trading rules and find that, while they are not generally optimal, they may perform better than fundamental analysis when traders are uncertain about the correct model of asset payoffs. However, both papers rely on highly specified models, and they do not directly address the relationship between uncertainty and

technical analysis. To explain technical trading as the result of a learning process or a state of mind, is to replace an observable anomaly with an unobservable phenomenon. Such an explanation can offer few testable predictions. On the other hand, if technical trading is a response to a specific informational problem, which might be analyzed or measured directly, then we should be able to predict when this behavior will occur.

2.2 Knightian Uncertainty

Non-probabilistic (or ‘Knightian’) uncertainty played a major role in the work of two pre-war economists, Frank Knight and John Maynard Keynes. Knight (1921) defined uncertainty as “the degree of subjective confidence felt in [a] contemplated act as a correct adaptation to the future” (p.242), and suggested that entrepreneurs earn economic profits as a reward for contending with and specializing in “the uncertainty which is involved in a changing world” (p.147). Similarly, Keynes (1936) argued that “investment based on genuine long-term expectation is so difficult today as to be scarcely practicable” (12.v), and that investment decisions consequently depend upon confidence, or “how highly we rate the likelihood of our best forecast turning out quite wrong” (12.ii). The common insight appears to have been that, while we might construct models of the economy that assign probabilities to events, we do not typically possess probabilistic knowledge about whether these models are correct. Market behavior may depend upon the amount of ‘confidence’ attached to market forecasts, and if forecasting models are judged based upon their empirical performance - as is typically done in the sciences - then the state of confidence may vary over time and in response to changes in economic data.

Neither economist developed uncertainty into a tractable theory. Knight’s analysis was “lacking in formal clarity” (Arrow, 1951), whereas Keynes believed that there is “not much to be said about the state of confidence a priori” (Keynes, 1936 12.ii). Moreover, it was difficult to reconcile uncertainty with the theoretical traditions of equilibrium analysis and expected utility, which assume that behavior is fixed and that outcomes have objective probabilities. The concept of uncertainty nevertheless played an important role in post-war macroeconomics. Macroeconomic models of the 1950’s and 60’s tended to be heavily statistical in nature, reflecting economists’ lack of confidence in their mechanical understanding of the broader economy, and agents within these models were also assumed to have statistical or ‘adaptive’ expectations. The adaptive expectations approach represented a “clear-cut conflict” (Lucas, 1976) with economic theory, and one that was eventually resolved with the

widespread adoption of rational expectations theory. Rational expectations imposes the restriction that, within an economic model, “the conditional expectation operator is consistent with the probability law that governs the actual data generation” (Hansen, 2014). In other words, agents within a probabilistic model have probabilistic expectations. While internally consistent, this approach rules out uncertainty, and it implies that agents do not face the same informational constraints - the same issue of confidence - as the economist, policymaker or market participant who built the model.

Although adaptive expectations was found to be a poor macroeconomic framework, it may nevertheless provide what Knight and Keynes did not: a way to model uncertainty in theoretical settings. This possibility has received considerable attention in recent years. Evans and Honkapohja (2001) propose an adaptive learning approach to modeling, in which agents are not presumed to have “full knowledge of the structure of the model” or “the values of the parameters,” but rather to “act like statisticians or econometricians when ... forecasting about the future state of the economy” (p.12). Cogley and Sargent (2005) posit that central banks are “adaptive decision makers” subject to model uncertainty, and that economic data not only informs each of the models that central bankers use, but influences as well the confidence assigned to the different models. Hansen (2014) proposes an empirical approach to asset pricing in which “a researcher does not specify formally [the probability law that governs the actual data generation] and instead ‘lets the data speak’.” Similarly, Frydman and Goldberg (2013) offer a model of asset prices in which some variables are left unspecified, to reflect the fact that “participants revise their forecasting strategies” as news becomes available.

Why does any of this matter with respect to technical trading? The broad answer is that, given Knightian uncertainty, prices may convey information about the forecasting models used by market participants. If these models (or their parameters) are chosen based upon empirical performance, then they will vary in response to changes in asset payoffs, and this learning process may introduce persistent trends into price movements. In other words, technical analysis may allow one to “guess better than the crowd how the crowd will behave” (Keynes, 1936 12.v). This intuition has some grounding in the literature. Zhang (2006) finds that short-term price momentum is correlated with informational uncertainty, or ambiguity regarding “the implications of new information for a firm’s value,” suggesting that learning may produce momentum in prices. Lawson (1980) shows that, given adaptive expectations

about inflation, “higher order” learning rules can be optimal when the inflation rate changes in non-constant fashion. A higher-order learning rule means, essentially, that agents are learning about the learning process. Returning to a point made in the previous section, uncertainty can potentially explain the stylized facts offered by Menkhoff and Taylor (2007): the correlation between technical profits and volatility, the time-varying nature of technical rule performance, and their greater use at short time horizons. Finally, there exists a sizable literature on the measurement of uncertainty using variation in economic indicators (Jurado et al, 2013) and the dispersion of analyst forecasts (Anderson et al, 2009; Barron et al, 2010), suggesting that a connection between uncertainty and technical analysis can be empirically tested.

3 Model

In this section I construct an asset-pricing model with the goal of testing, in a theoretical setting, the hypothesis that technical analysis is predictive under Knightian uncertainty. An asset-pricing model usually includes two elements: a stochastic process describing asset payoffs, and a ‘pricing kernel’ that translates future payoffs into present prices. To introduce uncertainty into such a model, I make two critical assumptions. First, I assume that agents do not know the parameters of the underlying stochastic process, but must form statistical estimates of these parameters. Second, in order to ensure that uncertainty is persistent, I assume that the model parameters change over time, and that agents remain ignorant of the “law of the change” (Knight, 1921 p.37). In other words, agents are not aware of all of the variables affecting asset payoffs, and they possess only imperfect information about those variables that they are aware of.

The model is constructed as follows. First, I define a stochastic process representing the dividends paid by a risky asset. I then define three prices: the ‘true price’ indicated by the stochastic process, the ‘fundamental price’ indicated by agents’ expectations about the stochastic process, and the ‘technical price’ indicated by agents’ expectations about prices. Since the fundamental and technical prices are both estimates of the true price, we can measure the relative accuracy of the technical price, and we can test whether its performance correlates with changes in the model parameters. Thus, this model may allow us to test the hypothesis that technical analysis is predictive under Knightian uncertainty.

3.1 Stochastic Process

In order to construct the model outlined above, we require a stochastic process with time-varying parameters. One of the simplest models satisfying this description is a Markov-switching stochastic process, in which one (or several) of the parameters in a simple linear model is made to depend upon an underlying ‘state’ variable. The state variable is assumed to follow a Markov process, meaning that its value in any given period depends only upon its value in the preceding period - an assumption that allows us to control the persistence of each state. Studied in-depth by Hamilton (1989), Markov-switching processes are a relatively simple way to characterize time series data that move in repetitive, well-defined trends, such as data influenced by the business cycle. These models have been used successfully to describe macroeconomic variables such as GDP (Kim et al, 2005) and equity returns (Guidolin and Timmerman, 2005).

I assume that there exists an asset whose dividends are given by a Markov-switching process with two states: an expansionary state (state one), and a contractionary state (state two). This stochastic process is described by the following set of equations:

$$\begin{aligned} \frac{Y_t}{Y_{t-1}} &= a_{s_t} + \epsilon_t \quad , \quad \epsilon_t \sim N(0, \sigma^2) \\ \mathbf{P} &= \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix} \\ p_{ij} &= P(s_t = j \mid s_{t-1} = i), \end{aligned} \tag{1}$$

where Y_t/Y_{t-1} is the slope of the dividend at period t , and $a_{s_t} \in \{a_1, a_2\}$ is a trend coefficient conditional on the state variable $s_t \in \{s_1, s_2\}$. The transition probabilities for the state variable are given by the matrix \mathbf{P} , where $p_{ij} \in [0, 1]$ indicates the probability that $s_t = s_j$ given that $s_{t-1} = s_i$. The random variable or ‘innovation’ ϵ_t is assumed to be independent and identically distributed. These equations state that the period-to-period change in asset dividends is determined by two components: a random component, and a trend component whose value varies over time.

Hamilton (1989) provides a probabilistic filter for estimating the value of the state variable in any period, given knowledge of the model’s parameters and information on all past dividends. Knightian uncertainty precludes agents from using this filter, but it is required for the model estimation performed in section 4, and so I describe it here. Hamilton gives

three equations:

$$\hat{s}_{j,t|t-1} = \sum_{i=1}^2 p_{ij} s_{i,t|t-1} \quad (2)$$

$$\eta_{j,t} = \left(\frac{1}{\sqrt{2\pi\sigma^2 Y_{t-1}^2}} \right) e^{-\frac{(Y_t - a_j Y_{t-1})^2}{2\sigma^2}} \quad (3)$$

$$\hat{s}_{j,t} = \frac{\sum_{i=1}^2 p_{ij} s_{i,t|t-1} \eta_{j,t}}{\sum_{i=1}^2 \sum_{j=1}^2 p_{ij} s_{i,t|t-1} \eta_{j,t}}, \quad (4)$$

where (2) is the ‘predicting’ equation, or the estimated probability that $s_t = j$ given dividend values up to period $t-1$, and (4) is the ‘updating’ equation, or the estimate of this probability using dividend values up to period t . In equation (3), $\eta_{j,t}$ indicates the probability density of dividend values at time t , conditional on $s_t = j$. Model estimation is discussed in greater detail in section 4.

3.2 Fundamental Rule

For reasons discussed at the beginning of this section, I assume that agents are ignorant of the ‘state’ variable described above. They form expectations about future asset dividends using the equation

$$\frac{Y_t}{Y_{t-1}} = a_t + \epsilon_t \quad , \quad \epsilon_t \sim N(0, \sigma^2) \quad , \quad (5)$$

where a_t varies over time and must be statistically estimated. For situations such as this, where agents form adaptive expectations about a stochastic process with time-varying parameters, Sargent (1999) suggests using a constant gain least squares (CGLS) rule, which gives greater weight to more recent observations. For the trend coefficient a_t in (9), a CGLS rule is given by the equations

$$E[a_t] = \hat{a}_t = \hat{a}_{t-1} + \lambda_F Y_{t-1} \frac{Y_t - \hat{a}_{t-1} Y_{t-1}}{R_t} \quad (6)$$

$$R_t = R_{t-1} + \lambda_F (Y_{t-1}^2 - R_{t-1}) \quad , \quad (7)$$

where $\lambda_F \in [0, 1]$ is the weight or ‘gain’ placed on the most recent observation. A higher gain means that the estimate \hat{a}_t will adjust more quickly to changes in the true parameter a_{s_t} , but that it will also be more volatile due to the randomness imposed by ϵ_t . I will refer to this estimate as the ‘fundamental rule.’

For derivation of price, I assume for simplicity that agents are risk neutral. I assume also that $E[Y_{t+1}] = \hat{a}_t Y_t$ and $E[\hat{a}_{t+1}] = \hat{a}_t$, or that agents form expectations using all of the information (and only the information) available to them. Finally, I assume that the price at time t is simply the present value of future dividends, with discount rate k . We therefore have the pricing equation:

$$\begin{aligned} P_t^F &= \frac{E[Y_{t+1}]}{1+k} + \frac{E[Y_{t+2}]}{(1+k)^2} + \dots + \frac{E[Y_{t+n}]}{(1+k)^n} + \dots \\ &= \sum_{i=1}^{\infty} \frac{E[Y_{t+i}]}{(1+k)^i}. \end{aligned} \quad (8)$$

Substituting (5) and (6) into this equation, we have

$$\begin{aligned} P_t^F &= \sum_{i=1}^{\infty} \left(\frac{\hat{a}_t}{1+k}\right)^i Y_t \\ &= \frac{Y_t}{1 - \frac{\hat{a}_t}{1+k}}, \end{aligned} \quad (9)$$

provided that $|\hat{a}_t| < (1+k)$. If this condition is not true, and $|\hat{a}_t| \geq (1+k)$, then the summation in (9) will not be a geometric series and it will yield an infinite or an undefined value. This is an important caveat, because it limits the values that we can assign to the discount rate k . I hereafter refer to P_t^F as the ‘fundamental price’.

3.3 Technical Rule

In the previous section I defined the fundamental rule and, from that rule, I derived the fundamental price. Note that this price will have a slope, equal in expected terms to the slope of dividends. That is,

$$E[P_{t+1}^F] = \frac{E[Y_{t+1}]}{1 - \frac{E[\hat{a}_{t+1}]}{1+k}} = \frac{\hat{a}_t Y_t}{1 - \frac{\hat{a}_t}{1+k}} = \hat{a}_t P_t^F. \quad (10)$$

Note additionally that price can be represented, like dividends, as a stochastic process of the form given in (3):

$$\frac{P_t^F}{P_{t-1}^F} = b_t + \theta_t, \quad \theta_t \sim N(0, \sigma_\theta^2). \quad (11)$$

Equation (10) shows that we could use the fundamental rule \hat{a}_t as an estimate of the slope of price, or b_t . However, we can also estimate b_t directly by using the CGLS equations given in (4) and (5), such that

$$E[b_t] = \hat{b}_t = \hat{b}_{t-1} + \lambda_T P_{t-1}^F \frac{P_t^F - \hat{b}_{t-1} P_{t-1}^F}{Q_t} \quad (12)$$

$$Q_t = Q_{t-1} + \lambda_T ([P_{t-1}^F]^2 - Q_{t-1}). \quad (13)$$

We therefore have a second adaptive expectation with gain λ_T , which treats price as a stochastic process and forms estimates about the period-to-period change in prices. I will refer to this expectation as the ‘technical rule.’ Note that $E[P_{t+1}^F] = \hat{a}_t P_t^F = \hat{b}_t P_t^F$, showing that each of the CGLS rules gives the expected slope of price. However, the estimates given by these two rules will differ, insofar as \hat{b}_t will depend upon variation in \hat{a}_t , as well as upon variation in the underlying dividend. This can be seen by substituting (9) into (12):

$$\hat{b}_t = \hat{b}_{t-1} + \frac{\lambda_T}{Q_t} \left(\frac{Y_t Y_{t-1} (1+k)^2}{(1+k - \hat{a}_{t-1})(1+k - \hat{a}_t)} - \frac{\hat{b}_t Y_{t-1}^2 (1+k)^2}{(1+k - \hat{a}_{t-1})^2} \right). \quad (14)$$

If \hat{a}_t is not equal to \hat{a}_{t-1} , then in general \hat{b}_t will not be equal to \hat{a}_t . When the two rules have equal gains, \hat{b}_t will tend to vary more than \hat{a}_t . Whether this increased variation brings with it increased predictive power will depend upon whether changes in \hat{a}_t are due to randomness, or due to a change in the model parameters - that is, to a state shift.

From \hat{b}_t we can derive a second, technical price:

$$P_t^T = \frac{Y_t}{1 - \frac{\hat{b}_t}{1+k}}. \quad (15)$$

I hypothesize that, following a state shift, the technical price will be a more accurate indicator than the fundamental price of the true present value of future dividends. To compare the two prices, I will make reference to a third price, given not by the estimates \hat{a}_t or \hat{b}_t but by the true parameter value a_{s_t} :

$$P_t^* = \frac{Y_t}{1 - \frac{a_{s_t}}{1+k}}. \quad (16)$$

This ‘true price’ provides a benchmark by which we might measure the accuracy of the fundamental and technical rules as predictors of future dividends, thereby determining the relative performance of the technical rule.

4 Results

To test the model just described, we require parameter values for the stochastic process $\{a_1, a_2, \sigma, P_{11}, P_{22}\}$, the discount rate k , and the gains of the two CGLS rules $\{\hat{\lambda}_F, \hat{\lambda}_T\}$. The approach I take in this section is as follows. I estimate the stochastic process using S&P 500 monthly Earnings-Per-Share (EPS) data provided by Shiller (2015), as well as quarterly EPS for four S&P 500 components: Microsoft, Exxon, General Electric and Johnson & Johnson.

Estimation is performed by maximizing the model log-likelihood with the Broyden-Fletcher-Goldfarb-Shanno algorithm, a numerical approximation method; the model log-likelihood is the summation over all periods of the log of the denominator in (4). Using these estimates, I examine technical rule performance for different values of the fundamental and technical gain. Specifically, I run a 100×100 grid of model simulations, each with $N = 5000$ periods, in which the fundamental and technical gains range from 0 to 1 in .01 increments. The performance of the technical rule is measured as the mean absolute percentage error (MAPE) of the fundamental price P_t^F as a predictor of the true price P_t^* , minus the MAPE of the technical price P_t^T :

$$MAPE_F - MAPE_T = \sum_{i=1}^N \frac{|P_i^F - P_i^*| - |P_i^T - P_i^*|}{N \times P_i^*}. \quad (17)$$

I report two general findings. First, for a significant range of gain values, the technical rule outperforms the fundamental rule in the periods following a shift in the state variable. Second, this outperformance is persistent over time and is robust to changes in parameter values, suggesting that the technical rule could be employed profitably by the agents in this model. However, these results are affected by the discount rate k , a free parameter. In order to address this shortcoming, in section 4.3 I characterize results under alternative discount rates.

4.1 S&P 500 Monthly Earnings

Here I estimate the model presented in section 3 using CPI-adjusted S&P 500 monthly EPS data provided by Shiller (2015). Data span the period from January 1871 to March 2015. However, a problem encountered during estimation is that the large decline in earnings for the S&P 500 during the 2007-08 financial crisis tends to dominate the time series, so that instead of trends corresponding to economic expansion and recession, we obtain trends corresponding to normal market conditions and financial crises. Since my goal in using index-wide earnings data is to examine the ‘typical’ performance of the technical rule, I wish to avoid any variation due to the financial crisis. Consequently, I limit model estimation to earnings data from the twentieth century (January 1901 to December 2000). Technical rule performance under more extreme conditions is addressed in section 4.2, where I estimate the model using individual corporate EPS.

Table 1: Parameter estimates for S&P 500 monthly EPS

Parameter	Estimate	Std. Error
a_1	1.015	.0012
a_2	0.985	.0013
σ	0.018	.0004
p_{11}	0.955	.0100
p_{22}	0.947	.0012

Compared to the null hypothesis of a linear (non-switching) model, the Markov-switching model from section 3 provides a good fit for monthly S&P 500 earnings. A likelihood ratio test, which compares the log likelihood values of the two models, yields a p-value $< .0001$. The performance of the technical rule is presented graphically in Figure 1, measured over three-month intervals beginning with the first month following a state variable shift. Darker areas indicate values of the fundamental and technical gain for which the technical price produces a lower MAPE than does the Fundamental MAPE - that is, where the technical rule has predictive power.

With the model estimates in table 1 and a discount rate of $k = .2$, the technical rule demonstrates predictive power for a wide range of fundamental and technical gain values. For ‘reasonable’ values of the fundamental gain (between .05 and .4)¹, technical rule performance is strongest in the months immediately following a state shift, supporting the hypothesis. For high values of the fundamental gain, this relationship is reversed; the fundamental rule adjusts quickly to state shifts, but performs worse over time than does a technical rule with a low gain. However, this second result requires both unreasonably high values of the fundamental gain, and a significant disparity between the fundamental and technical gains. For more reasonable gain values with less disparity, the hypothesis holds: technical rule performance correlates with Knightian uncertainty, as indicated by the recency of a state shift.

¹For the model estimates presented in this thesis, ‘optimal’ fundamental gains range from .17 to .36, with standard errors no greater than .009. Although Knightian uncertainty rules out such optimization, these figures nevertheless indicate a ‘reasonable’ range of gain values.

Figure 1: Technical Rule Performance, S&P 500 Monthly EPS

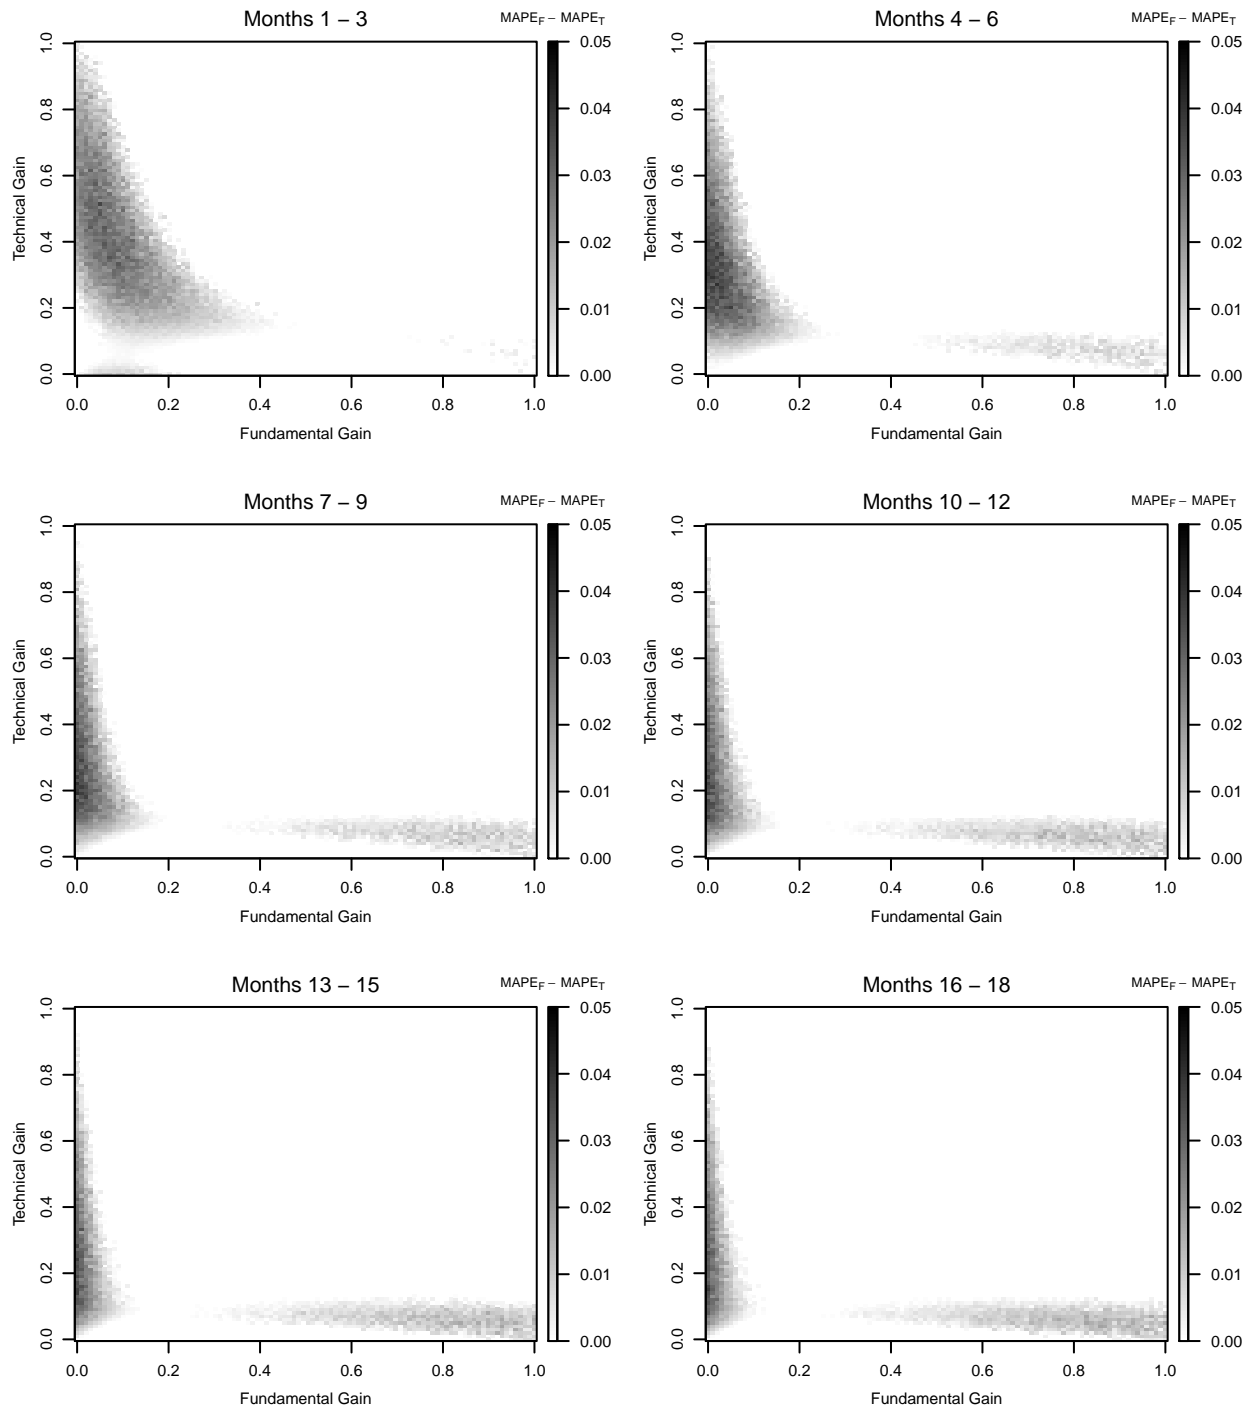


Figure 1: Technical rule performance in the months following a state shift, using S&P 500 monthly earnings. Darker areas indicate gain values where the technical rule outperforms the fundamental rule.

4.2 Corporate Quarterly Earnings

To examine how different model parameters affect technical rule performance, in this section I estimate the Markov-switching model using quarterly EPS for four of the five² largest components of the S&P 500 stock index: Microsoft, Exxon, General Electric and Johnson and Johnson. I obtain data from Compustat, beginning in Q1 1962 for Exxon and GE, Q1 1963 for Johnson and Johnson, and Q3 1984 for Microsoft. All time series end in Q2 2015. I present this data graphically in figure 2.




Figure 2: Log EPS for Exxon, GE, J&J and Microsoft

Because the Hamilton filter described in section 3.1 cannot handle outliers or large ‘jumps’ - they introduce division by zero into equation (4) - the data required smoothing before model estimates could be obtained. I found that a four-quarter moving average did not sufficiently eliminate breaks in the data, so instead I used the annual growth rate of earnings, measured quarterly. Model estimates are presented in table 2.

Table 2: Parameter estimates for corporate quarterly EPS

	Exxon		GE		J&J		Microsoft	
Parameter	Estimate	Error	Estimate	Error	Estimate	Error	Estimate	Error
a_1	1.084	.012	1.016	.003	1.029	.003	1.081	.005
a_2	0.986	.006	0.834	.018	0.887	.012	0.967	.009
σ	0.047	.003	0.043	.002	0.040	.002	0.038	.003
p_{11}	0.839	.054	0.985	.009	0.984	.009	0.930	.028
p_{22}	0.945	.012	0.771	.003	0.785	.003	0.834	.005
p-value	< .0001		< .0001		< .0001		< .0001	

²The volatility of Apple earnings prevented me from obtaining a model estimate.

Likelihood ratio tests yield a p-value $< .0001$ in all four cases, when compared against the null hypothesis of a linear model. The EPS of each company exhibits an expansionary and a contractionary trend. For Microsoft and Exxon, the difference between these trends is relatively mild, and is indicative of the kind of business cycle variation discussed in section 4.1. Note that in Exxon's case, the expansionary trend is less persistent than the expansionary trend, reflecting the characteristics of the oil industry. On the other hand, earnings for GE and Johnson & Johnson are characterized by a more severe, but less persistent contractionary trend, reflecting the influence of the financial crisis on GE's performance, and of product recalls (as in 2010) on that of J&J.

Figure 3: Technical Rule Performance, Corporate Quarterly EPS




Figure 3: Technical rule performance in the four quarters following a state shift, using corporate EPS data. Darker areas indicate gain values where the technical rule outperforms the fundamental rule. Note the difference in scales.

Figure 3 shows that the technical rule has considerably greater predictive power in the case of GE and J&J. Also visible is the effect of the larger standard deviation (σ) of Exxon earnings relative to Microsoft earnings, and the narrower spread between trends for J&J when compared to GE. Overall, technical rule performance tends to improve as the standard deviation of earnings decreases, or as the difference between two trends widens - in other words, as state shifts account for more of the variation in earnings.

4.3 Alternative Discount Rates

The simulations in previous sections use the discount rate $k = .2$. This value was chosen to be sufficiently high that the pricing equations described in section 3 would yield well-behaved (i.e. finite) values. Although the fundamental rule \hat{a}_t can have a value no greater than the larger trend coefficient plus the innovation variable, and no less than the smaller trend coefficient minus the innovation, the technical rule can potentially be much larger since it also takes into account variation in \hat{a}_t . Because there exists a vertical asymptote in the technical price as \hat{b}_t approaches k , this nonlinearity can affect the MAPE of the technical price when compared to the fundamental price. In order to characterize the effect of a change in the discount rate on technical rule performance, I simulate S&P 500 monthly earnings for months 1-3 and 4-6 as in section 4.1, but instead I use the discount values $k = .1$ and $k = .3$. Results are presented in figure 3.

All else equal, a higher discount rate tends to flatten results, giving the technical rule predictive power over a wider range of gain values, but reducing the magnitude of that predictive power. Note that for a lower discount rate, a low-gain technical rule no longer outperforms a high-gain fundamental rule. Although I argued in section 4.1 that results in this portion of the graph require unreasonable gain values, the explanation for this change is nevertheless informative. As the discount rate becomes smaller, the impact of a state shift on price becomes relatively greater, resulting in a general increase in the MAPE of a technical rule with a low gain, relative to the MAPE of a fundamental rule with a high gain.

Overall, a change in the discount rate does not significantly affect previous results. Rather, it reflects a tradeoff: a higher discount rate means that the technical rule is less likely to be predictive, but that it will have a larger advantage over the fundamental rule in cases where it is predictive.

Figure 4: Alternative Discount Rates, S&P 500 Monthly EPS

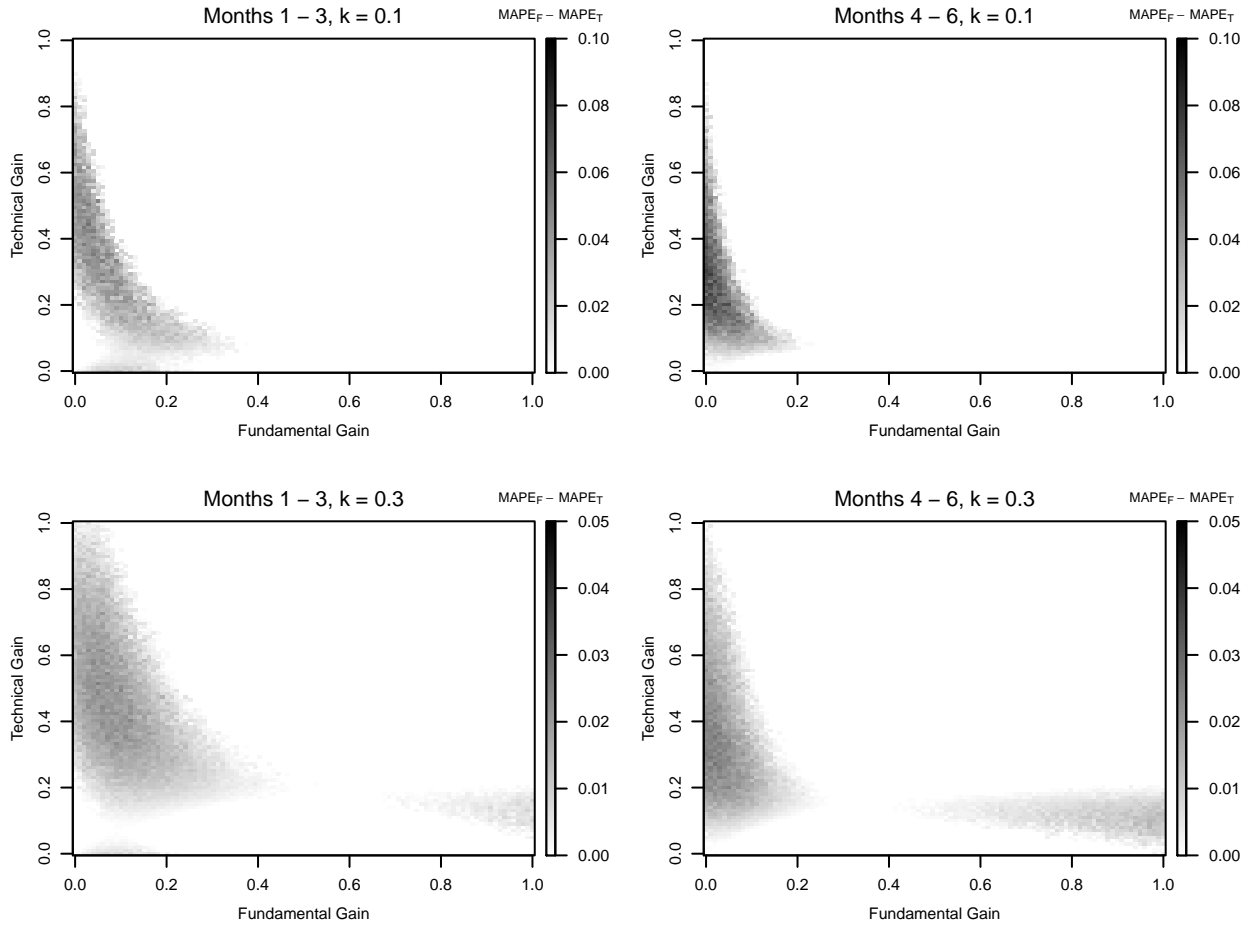


Figure 4: Technical rule performance in the months following a state shift, using S&P 500 monthly earnings. Darker areas indicate gain values where the technical rule outperforms the fundamental rule. Note the difference in scales.

5 Conclusion

The results in section 4 show that, for a wide range of parameter values, the technical rule has predictive power in the periods following a state shift. This predictive power is persistent, suggesting that the gains offered by the technical rule are exploitable. The relationship between model parameters and technical rule performance can be characterized as follows:

- i) For plausible values of the fundamental gain, technical rule performance is strongest in the periods immediately following a state shift.
- ii) The predictive power of the technical rule tends to be more persistent over time for low

values of the fundamental and technical gains - that is, when agents learn slowly.

- iii) A decrease in the standard deviation, and/or an increase in the absolute difference between the two trend coefficients, tends to increase the predictive power of the technical rule relative to the fundamental rule.
- iv) An increase in the discount rate tends to expand the range of gain values for which the technical rule is predictive, but reduces the predictive power of the technical rule within that range.

Generalizing these results, we can say that technical rule performance improves as the Markov-switching process becomes more influential - that is, as Knightian uncertainty increases. These results therefore lend support to the hypothesis that, given Knightian uncertainty, a 'rational' investor might choose to employ technical analysis.

Some caveats are in order when interpreting this evidence. First, while it is common in economics to test hypotheses in theoretical models, such models can only assess whether a hypothesis is internally consistent and theoretically plausible. Empirical research is needed to determine whether the model predictions are consistent with economic data. Thus, while this thesis represents a first step in testing the relationship between technical analysis and uncertainty, it represents a first step only. Second, I make strong assumptions about the informational constraints facing agents: that they have neither perfect information about, nor a perfect theoretical understanding of asset payoffs. While I argue that these assumptions are realistic, they are nevertheless at odds with rational expectations theory. Third, simulation results show that technical rule performance is conditional upon the value of the fundamental and technical gains. Even if we make the limiting assumption that all agents learn at the same rate, or $\lambda_F = \lambda_T$, there will still exist a large range of gain values for which technical analysis is not predictive. Thus, while results indicate that technical analysis can be useful, they do not require it. Fourth, technical rule performance may depend upon the choice of CGLS for the model in section 3, as opposed to some other statistical method for estimating model parameters. Finally, the fundamental price is an input for the technical rule, but this price could only be observable if no traders were using the technical rule. In other words, the use of technical analysis may adversely affect its performance. The goal of this thesis was to investigate whether technical analysis is a reasonable behavior *a priori*, but this leaves open the important question of what happens once technical traders are present.

If technical analysis is predictive under Knightian uncertainty, then several implications follow. From the point of view of investors, there may be circumstances where it is normative to take technical considerations into account - contradicting the advice offered by mainstream asset-pricing theories. From the standpoint of financial institutions, who typically have regulatory and fiduciary duties to monitor their risk exposure, a better understanding of the role of uncertainty in driving ‘anomalous’ market behavior may allow these institutions to better hedge or protect themselves against extreme market conditions. Similarly, knowledge of any relationship between technical analysis and uncertainty may assist policymakers in designing efficient policy, and in overseeing the operations of financial institutions. With respect to economic theory, if it can be shown that uncertainty is responsible for a major behavioral phenomenon, then there may be a greater impetus to study the kind of informational constraints discussed here, and to consider them when building economic models.

In conclusion, this thesis provides a framework for studying the relationship between uncertainty and technical analysis, and theoretical evidence in favor of the hypothesis that such a relationship exists. The next step would be to test whether this relationship is empirically supported. Does short-term asset price momentum correlate with higher variance in analyst forecasts? Does price volatility increase when economic data is subject to greater noise? The results presented here suggest that these questions may be worth answering. More broadly, this thesis is an investigation into how we might explain or even predict anomalous behavior in asset markets, by considering the informational constraints facing market participants and the uncertainty caused by these constraints. Understanding such behavior allows us to better assess market risk, to better design robust institutions and policy, and to better reconcile theory with observation. Ultimately, asset traders and economic models face the same challenge: surviving unforeseen events and market crises. With increased knowledge of the relationship between uncertainty and asset prices, market participants and economic forecasters may be better able to meet that challenge.

References

- [1] Arrow, Kenneth (1951). “Alternative Approaches to the Theory of Choice in Risk-Taking Situations.” *Econometrica* 19.4: 404-437.
- [2] Bittlingmayer, George (1998). “Output, Stock Volatility, and Political Uncertainty in a Natural Experiment: Germany, 1880-1940.” *Journal of Finance* 53.6: 2243-2257.
- [3] Brock, William and Josef Lakonishok and Blake LeBaron (1992). “Simple Technical Trading Rules and the Stochastic Properties of Stock Returns.” *Journal of Finance* 47: 1731-1764.
- [4] Brogaard, Jonathan and Terrence Hendershott and Ryan Riordan (2014). “High-Frequency Trading and Price Discovery.” *Review of Financial Studies* 27.8: 2267-2306.
- [5] Brorsen, Wade and Scott Irwin (1987). “Futures Funds and Price Volatility.” *The Review of Futures Markets* 6: 118-135.
- [6] Cogley, Timothy and Thomas Sargent (2005). “The Conquest of US Inflation: Learning and Robustness to Model Uncertainty.” *Review of Economic Dynamics* 8.2: 528-563.
- [7] De Long, Bradford and Andrei Schleifer, Lawrence Summers and Robert Waldmann (1990). “Noise Trader Risk in Financial Markets.” *Journal of Political Economy* 98.4: 703-738.
- [8] Evans, George and Seppo Honkapohja (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton.
- [9] Fama, Eugene and Marshall Blume (1966). “Filter Rules and Stock-Market Trading.” *Journal of Business* 39: 226-241.
- [10] Frydman, Roman and Michael Goldberg (2013). “Opening Models of Asset Prices and Risk to Non-Routine Change.” *Rethinking Expectations: The Way Forward For Macroeconomics*. Ed. R. Frydman and E. Phelps. Princeton Press, Princeton: 207-250.
- [11] Guidolin, Massimo and Allan Timmermann (2005). “Economic Implications of Bull and Bear Regimes in UK Stock and Bond Returns.” *The Economic Journal* 115.500: 111-143.
- [12] Hamilton, James (1989). “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle.” *Econometrica* 57.2: 357-384.

- [13] Hansen, Lars (2014). “Uncertainty Outside and Inside Economic Models.” *National Bureau of Economic Research* No. w20394.
- [14] Joshi, Shareen and Jeffrey Parker and Mark Bedau (2000). “Technical Trading Creates a Prisoner’s Dilemma: Results From an Agent-Based Model.” *Computational Finance* 99: 465-479.
- [15] Keynes, John Maynard (1936). *General Theory of Employment, Interest and Money*. Palgrave, Macmillan, London.
- [16] Kim, Chang-Jin and James Morley and Jeremy Piger (2005). “Nonlinearity and the Permanent Effects of Recessions.” *Journal of Applied Econometrics* 20.2: 291-309.
- [17] Knight, Frank (1921). *Risk, Uncertainty and Profit*. Houghton Mifflin, Boston.
- [18] Lawson, Tony (1980). “Adaptive Expectations and Uncertainty.” *The Review of Economic Studies*: 305-320.
- [19] Lucas, Robert (1976). “Econometric Policy Evaluation: A Critique.” *Carnegie-Rochester Conference Series on Public Policy* 1: 19-46.
- [20] Lux, Thomas and Michele Marchesi (1998). “Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market.” *Nature* 397.6719: 498-500.
- [21] Malkiel, Burton (1973). *A Random Walk Down Wall Street*. Norton & Company, New York.
- [22] Manzoor, Sarfraz (2013). “Quants: the Maths Geniuses Running Wall Street.” *The Telegraph*, 23 July.
- [23] Menkhoff, Lukas (2010). “The Use of Technical Analysis by Fund Managers: International Evidence.” *Journal of Banking & Finance* 34.11: 2573-2586
- [24] Menkhoff, Lukas and Mark Taylor (2007). “The Obstinate Passion of Foreign Exchange Professionals: Technical Analysis.” *Journal of Economic Literature*: 936-972.
- [25] Neely, Christopher and Paul Weller and Rob Dittmar (1997). “Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach.” *Journal of Financial and Quantitative Analysis* 32.4: 405-426.

- [26] Oberlechner, Thomas (2001). “Importance of Technical and Fundamental Analysis in the European Foreign Exchange Market.” *International Journal of Finance and Economics* 6.1: 91-93.
- [27] Park, Cheol-Ho and Scott Irwin (2007). “What Do We Know About the Profitability of Technical Analysis?” *Journal of Economic Surveys* 21.4: 786-826.
- [28] Poterba, James, and Lawrence Summers (1988). “Mean Reversion in Stock Prices: Evidence and Implications.” *Journal of Financial Economics* 22.1: 27-59.
- [29] Pring, Martin (2002). *Technical Analysis Explained*. McGraw-Hill, New York.
- [30] Shiller, Robert (1981). “Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?” *The American Economic Review* 71.3: 421-436.
- [31] Shiller, Robert (2015). *Irrational exuberance*. Princeton University Press, Princeton.
- [32] Smidt, Seymour (1965). *Amateur Speculators*. Cornell Studies in Policy and Administration. Cornell University, New York.
- [33] Sweeney, Richard (1986). “Beating the Foreign Exchange Market.” *Journal of Finance* 41: 163-182.
- [34] Taylor, Mark and Helen Allen (1992). “The Use of Technical Analysis In the Foreign Exchange Market.” *Journal of International Money and Finance* 11.3: 304-314.
- [35] Taylor, Stephen (2000). “Stock Index and Price Dynamics in the UK and the US: New Evidence from a Trading Rule and Statistical Analysis.’ *European Journal of Finance* 6: 39-69.
- [36] Zhu, Yingzi and Guofu Zhou (2009). “Technical Analysis: An Asset Allocation Perspective On the Use of Moving Averages.” *Journal of Financial Economics* 92.3: 519-544.

Appendix: R Code

```
# This function compares the performance of a technical rule and a fundamental rule after
# a regime switch, by solving for the difference between their mean average percentage error
# for all gain values in (0,1) in .01 increments. Results are exported to a file "Data.csv".
# The function takes as input the length of the random walk, and the first and last periods
# of the region after a regime shift over which to average results. Setting primer = "on"
# avoids coordinates that yielded extreme values in the previous trial, reducing computation
# time. Note that the first 100 periods are cut off when computing results.
```

```
library("plot3D")
```

```
fundtech <- function(walklength, testbegin, testend, primer = "off") {
  # Model Parameters
  rfrate <- .20 # Risk-free interest rate; low values may cause erratic price behavior
  stdev <- .018 # Innovation standard deviations as a fraction of value
  # Regime Parameters
  pmatrix <- matrix(c(.955, .053, .045, .947), 2, 2) # Transition matrix
  regimes <- c(1, 2) # Regime index
  divtrend <- c(1.015, .985) # Dividend drift as a fraction of value
  # Initial Conditions
  initialseed <- 1 # Starting random seed
  initialregime <- 1 # Starting regime
  initialvalue <- 1 # Starting dividend
  initialexpdivtrend <- sum(divtrend)/2
  initialexppricetrend <- sum(divtrend)/2

  # Execution Loop
  cutoff <- -.2
  results <- matrix(-.05, nrow = 100, ncol = 100)
  if(primer == "on") results <- as.matrix(read.table("Data.csv"))
  for(fundgain in 1:100) for(techgain in 1:100) if(results[fundgain,techgain] > cutoff) {
    regime <- initialregime
    set.seed(initialseed)

    # Generate Random Vectors
    innovation <- rnorm(walklength - 1)
    switch <- sample(1:1000, walklength - 1, replace = TRUE)
    for(i in 2:walklength) {
      for(j in 1:2) {
        if(regime[i-1] == regimes[j]) if(switch[i-1] <= pmatrix[j, j]*
          1000) regime[i] <- regimes[j] else regime[i] <-
```

```

        regimes[regimes != j]
    }
}

# Generate Value and True Price Vectors
value <- initialvalue
trueprice <- vector()
for(i in 2:walklength) value[i] <- divtrend[regime[i]]*value[i-1] +
    stdev*value[i-1]*innovation[i-1]
for(i in 1:walklength) trueprice[i] <- value[i]*(1 + rfrate)/(1 +
    rfrate - divtrend[regime[i]])

# Fundamental Price Vector
# Least Squares Equations
expdivtrend <- vector(); expdivtrend[1] <- initialexpdivtrend
expdivvar <- vector(); expdivvar[1] <- 1
for(i in 2:walklength){
    expdivvar[i] <- expdivvar[i-1] + (fundgain/100)*(value[i-1]^2 -
        expdivvar[i-1])
    expdivtrend[i] <- expdivtrend[i-1] + (fundgain/100)*value[i-1]*(value[i] -
        value[i-1]*expdivtrend[i-1])/expdivvar[i]
}

# Fundamental Price Vector
fundamentalprice <- vector()
for(i in 1:walklength) fundamentalprice[i] <- value[i]*(1 +
    rfrate)/(1 + rfrate - expdivtrend[i])

# Technical Price
# Least Squares Equations
exppricetrend <- vector(); exppricetrend[1] <- initialexppricetrend
exppricevar <- vector(); exppricevar[1] <- 1
for(i in 2:walklength){
    exppricevar[i] <- exppricevar[i-1] + (techgain/100)*
        (fundamentalprice[i-1]^2 - exppricevar[i-1])
    exppricetrend[i] <- exppricetrend[i-1] + (techgain/100)*
        fundamentalprice[i-1]*(fundamentalprice[i] -
        fundamentalprice[i-1]*exppricetrend[i-1])/exppricevar[i]
}

# Technical Price Vector
technicalprice <- vector()
for(i in 1:walklength) technicalprice[i] <- value[i]*(1 + rfrate)/(1 + rfrate -

```

```

    exppricetrend[i])

# Generate Mean Absolute Percentage Error Matrix
switchperiods <- vector() # Periods in which a regime shift occurs
for(i in 101:(walklength-testend)) for(j in 1:2) {
  if(regime[i-1] == regimes[j]) if(switch[i-1] > pmatrix[j, j]*
    1000) switchperiods = c(switchperiods, i)
}
mape <- matrix(0, nrow = 2, ncol = length(switchperiods)) # MAPE matrix
for(i in 1:length(switchperiods)) {
  for(j in 1:(1+testend-testbegin)) {
    mape[1,i] <- mape[1,i] + abs((fundamentalprice[switchperiods[i]-2+j+
      testbegin] - trueprice[switchperiods[i]-2+j+testbegin]))/
      (trueprice[switchperiods[i]-2+j+testbegin]*(1+testend-testbegin))
    mape[2,i] <- mape[2,i] + abs((technicalprice[switchperiods[i]-2+j+
      testbegin] - trueprice[switchperiods[i]-2+j+testbegin]))/
      (trueprice[switchperiods[i]-2+j+testbegin]*(1+testend-testbegin))
  }
}

# Generate test statistic, move random seed and print update
results[fundgain, techgain] <- mean(mape[1,]) - mean(mape[2,])
initialseed <- initialseed + 1
if(initialseed %% 100 == 0) print(paste(initialseed/100,"% complete.", sep = ""))
}

# Save Data File
write.table(results, file = "Data.csv")
}

# This function fits a two-regime Markov switching model to the data supplied, tests against
# the null hypothesis of a linear model using a likelihood ratio test, and reports the
# parameters of the two models, their log likelihoods, and the p-value for the likelihood
# ratio test.

mkestimate <- function(datavector){
  # Initial estimates <- c(coeff1, coeff2, stdev, p11, p22)
  initialparameters <- c(1.08, .98, .03, .98, .9) # Initial guess
  # Initial null estimates <- c(coeff, stdev)
  initialnullparameters <- c(1.01, .05)

```

```

# Fit the Markov switching models and return the log likelihood value
loglikelihood <- function(parameters) {
  pmatrix <- matrix(c(parameters[4], 1-parameters[5], 1-
    parameters[4], parameters[5]), 2, 2)
  # Hamilton Filter
  density <- array(0, dim = c(2, length(datavector)))
  jointconddensity <- array(0, dim = c(2, 2, length(datavector)))
  update <- array(0, dim = c(2, length(datavector))); update[,1] <- c(.5, .5)
  for(i in 2:length(datavector)) {
    for(k in 1:2) density[k, i] <- (1/(sqrt(2*pi)* abs(parameters[3]*
      datavector[i-1])))*exp(-(datavector[i] - parameters[k]*
      datavector[i-1])^2)/(2*(parameters[3]*datavector[i-1])^2))
    for(j in 1:2) for(k in 1:2) jointconddensity[j, k, i] <- pmatrix[j, k]*
      update[j, i-1]*density[k, i]
    for(k in 1:2) update[k, i] <- sum(jointconddensity[, k, i]) /
      sum(jointconddensity[, ,i])
  }
  loglikelihood <- 0
  for(i in 2:length(datavector)) loglikelihood <- loglikelihood +
    log(sum(jointconddensity[, ,i]))
  return(-loglikelihood)
}

# Fit the null model and return the log likelihood value
nullloglikelihood <- function(nullparameters) {
  parameters <- c(nullparameters[1],nullparameters[1],nullparameters[2],1,1)
  pmatrix <- matrix(c(parameters[4], 1-parameters[5], 1-
    parameters[4], parameters[5]), 2, 2)
  # Hamilton Filter
  density <- array(0, dim = c(2, length(datavector)))
  jointconddensity <- array(0, dim = c(2, 2, length(datavector)))
  update <- array(0, dim = c(2, length(datavector))); update[,1] <- c(.5, .5)
  for(i in 2:length(datavector)) {
    for(k in 1:2) density[k, i] <- (1/(sqrt(2*pi)*abs(parameters[3]*
      datavector[i-1])))*exp(-((datavector[i] - parameters[k]*
      datavector[i-1])^2)/(2*(parameters[3]*datavector[i-1])^2))
    for(j in 1:2) for(k in 1:2) jointconddensity[j, k, i] <- pmatrix[j, k]*
      update[j, i-1]*density[k, i]
    for(k in 1:2) update[k, i] <- sum(jointconddensity[, k, i]) /
      sum(jointconddensity[, ,i])
  }
}

```

```

    }
    loglikelihood <- 0
    for(i in 2:length(datavector)) loglikelihood <- loglikelihood +
      log(sum(jointcondensity[, ,i]))
    return(-loglikelihood) # Negative because optim() minimizes
  }

estimate <- optim(initialparameters, loglikelihood, method = "L-BFGS-B", hessian =
  TRUE, lower = c(.9,.5,.02,.01,.01), upper = c(1.5,1.2,.5,1,1)) # BFGS optimize
nullestimate <- optim(initialnullparameters, nullloglikelihood, method =
  "L-BFGS-B", hessian = TRUE, lower = c(.8, .02), upper = c(1.2, .2))
error <- sqrt(diag(solve(estimate$hessian))) # Standard errors from (neg) Hessian
nullerror <- sqrt(diag(solve(nullestimate$hessian)))
lratiosat <- 2*(nullestimate$value - estimate$value) # $value * (-1) = loglikelihood
lratiotest <- 1-pchisq(lratiosat, 3) # Calculate likelihood ratio test P-value

# Report
rpar <- round(estimate$par, 4); rnullpar <- round(nullestimate$par, 4)
rvalue <- round(-estimate$value, 4); rnullvalue <- round(-nullestimate$value, 4)
error <- round(error, 4); rnullerror <- round(nullerror, 4)
print(paste("Coefficient 1:", rpar[1], ", SE =", error[1], sep = " "))
print(paste("Coefficient 2:", rpar[2], ", SE =", error[2], sep = " "))
print(paste("Standard Deviation:", rpar[3], ", SE =", error[3], sep = " "))
print(paste("P11:", rpar[4], ", SE =", error[4], sep = " "))
print(paste("P22: ", rpar[5], ", SE =", error[1], sep = " "))
print(paste("Log likelihood value of", rvalue, "with convergence =",
  estimate$convergence, sep = " "))
print(paste("Null Coefficient:", rnullpar[1], ", SE =", rnullerror[1], sep = " "))
print(paste("Null Standard Deviation:", rnullpar[2], ", SE =", rnullerror[2], sep =
  " "))
print(paste("Null log likelihood value of", rnullvalue, "with convergence =",
  nullestimate$convergence, sep = " "))
print(paste("Likelihood Ratio Test P-value:", lratiotest, sep = " "))
}

# This function finds the optimal fundamental gain for the Markov-switching model
# using the BFGS algorithm.

mkoptimize <- function(walklength = 5100, seed = 1) {
  initialparameter <- .2 # Initial guess

```



```

mape <- function(fundgain) {
  # Model Parameters
  rfrate <- .20 # Risk-free interest rate; low values may produce N/A's
  stdev <- .018 # Innovation standard deviations as a fraction of value
  # Regime Parameters
  pmatrix <- matrix(c(.955, .053, .045, .947), 2, 2) # Transition matrix
  regimes <- c(1, 2) # Regime index
  divtrend <- c(1.015, .985) # Percent dividend trend under regimes 1 and 2
  # Initial Conditions
  set.seed(seed) # Random seed
  regime <- 1 # Starting regime
  initialvalue <- 1 # Starting dividend
  initialexpdivtrend <- sum(divtrend)/2
  initialexppricetrend <- sum(divtrend)/2

  # Generate Random Vectors
  innovation <- rnorm(walklength - 1)
  switch <- sample(1:1000, walklength - 1, replace = TRUE)
  for(i in 2:walklength) {
    for(j in 1:2) {
      if(regime[i-1] == regimes[j]) if(switch[i-1] <= pmatrix[j, j]*
        1000) regime[i] <- regimes[j] else regime[i] <-
        regimes[regimes != j]
    }
  }

  # Generate Value and True Price Vectors
  value <- initialvalue
  trueprice <- vector()
  for(i in 2:walklength) value[i] <- divtrend[regime[i]]*value[i-1] +
    stdev*value[i-1]*innovation[i-1]
  for(i in 1:walklength) trueprice[i] <- value[i]*(1 + rfrate)/(1 +
    rfrate - divtrend[regime[i]])

  # Least Squares Equations
  expdivtrend <- vector(); expdivtrend[1] <- initialexpdivtrend
  expdivvar <- vector(); expdivvar[1] <- 1
  for(i in 2:walklength){
    expdivvar[i] <- expdivvar[i-1] + (fundgain)*(value[i-1]^2 - expdivvar[i-1])
    expdivtrend[i] <- expdivtrend[i-1] + (fundgain)*value[i-1]*(value[i] -

```

```

        value[i-1]*expdivtrend[i-1])/expdivvar[i]
    }
    # Fundamental Price Vector
    fundamentalprice <- vector()
    for(i in 1:walklength) fundamentalprice[i] <- value[i]*(1 +
        rfrate)/(1 + rfrate - expdivtrend[i])

    # Mean Absolute Percentage Error
    pe <- 0
    for(i in 101:walklength) pe <- pe + abs((fundamentalprice[i] -
        trueprice[i]))/(trueprice[i]*(walklength-100))
    return(pe)
}

# Compute average optimized value with standard error
length <- 50
results <- rep(0, times = length)
for(i in 1:length) {
    estimate <- optim(initialparameter, mape, method = "L-BFGS-B", lower =
        .01, upper = 1, hessian = TRUE) # Optimization using BFGS)
    results[i] <- estimate$par
    seed <- seed + 1
}
avg <- mean(results)
stderr <- 0
for(i in 1:length) stderr <- stderr + ((results[i] - avg)^2)/(length-1)
stderr <- sqrt(stderr)/sqrt(length-1)

# Report
print(paste("Optimal Gain:", avg, sep = " "))
print(paste("Standard Error: ", stderr, sep = " "))
}

```