

Durham E-Theses

Three essays in behavioural finance: An examination into non-Bayesian Investment behaviour

ANTONIOU, CONSTANTINOS

How to cite:

ANTONIOU, CONSTANTINOS (2010) *Three essays in behavioural finance: An examination into non-Bayesian Investment behaviour*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/428/>

Use policy

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

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

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

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

Abstract

Thesis Title: Three essays in behavioural finance: An examination into non- Bayesian Investment behaviour

by

Constantinos Antoniou

“There is no general principle that prevents the creation of an economic theory based on other hypotheses than that of rationality”

Kenneth Arrow (1986)

Behavioural Finance relaxes the neoclassical assumption that investors consistently apply Bayes Rule when updating their expectations, and identifies the behavioural attributes that affect asset prices. This thesis extends this literature by examining deviations from the Bayesian model that arise due to *i*) ambiguity aversion, *ii*) investor sentiment and *iii*) decision heuristics.

Bayesian Updating assumes that investors are able to always estimate a *single* generating process for expected returns. However, in reality investors analyze noisy information signals that relate to this unknown distribution in a *latent* way, and it is likely that they are not *always* able to determine a *single* probability distribution. Behavioural economists have shown that in such conditions of uncertainty about probabilities people become pessimistic. The first chapter examines whether the pricing of analyst earnings is affected by ambiguity aversion, offering confirmatory evidence.

A behavioural literature shows that people in good sentiment make optimistic choices, relative to objective probabilities. The second chapter examines whether investor

sentiment affects the performance of the momentum trading strategy, an anomaly related to the pricing of good and bad information. The results indicate that sentiment strongly affects the momentum phenomenon, suggesting that it is triggered from investors' behavioural biases.

It has been suggested that deviations from Bayesian Updating arise due to heuristics triggered by the *characteristics* of the information used. The last chapter examines the validity of one such important hypothesis proposed by Griffin and Tversky (1992) using rigorous experimental economics techniques. The results confirm this hypothesis, indicating that investors are likely to overreact to salient information signals with low predictive validity.

Thesis Title:

Three Essays in Behavioural Finance:

An examination into non-Bayesian investment behaviour

by

Constantinos Antoniou

PhD in Economics

University of Durham

Durham Business School

September 2010

Table of Contents

Summary of contributions	10
1. Introduction.....	13
1.1 A shift of paradigm in financial economics	13
1.2 Behavioural Finance.....	15
1.3 Hypotheses and contributions	17
1.3.1 <i>Ambiguity in analyst forecast accuracy</i>	17
1.3.2 <i>The effect of Investors' Sentiment on Price Momentum</i>	19
1.3.3 <i>Decision Heuristics</i>	21
2. Bayes Rule and Investment decisions	24
2.1 Bayesian asset pricing	24
2.2 Bayesian asset pricing in practice	26
2.3 Behavioural asset pricing	27
2.3.1 <i>Heuristic simplification -Representativeness</i>	28
2.3.2 <i>Self deception- Overconfidence</i>	30
2.3.3 <i>Mood and Feelings- Investor sentiment</i>	33
2.3.4 <i>Non- Expected Utility Preferences</i>	37
2.4 Limits to Arbitrage	43
3. Ambiguity aversion and the pricing of analyst forecasts.....	47
3.1 Introduction	47
3.2 The theory of ambiguity aversion	52
3.3 Methodology and Data	54
3.3.1 <i>Definition of Variables and method that ambiguity is measured</i>	54
3.3.2 <i>The event study</i>	59
3.3.3 <i>Market reaction to ambiguity under general optimism and pessimism</i>	62
3.3.4 <i>Multivariate analysis</i>	63
3.3.5 <i>Sample construction and descriptive statistics</i>	66
3.4 Results	67
3.4.1 <i>Estimating forecast ambiguity</i>	67
3.4.2 <i>Impact and adjustment period returns by ambiguity proxy</i>	70
3.4.3 <i>Impact and Adjustment Period Returns and the Level of the Market</i>	73
3.4.4 <i>Multivariate analysis</i>	74
3.4.5 <i>Revision magnitude, company size and earnings announcements</i>	75
3.5 Further discussion	76

3.6 Conclusion.....	78
4. Investor Sentiment and Price Momentum.....	89
4.1 Introduction	89
4.2 Data and Methodology	94
4.3 The Empirical Evidence on Momentum Profits across Sentiment States	99
4.3.1 <i>Investor Sentiment and Short-Run Momentum Profits</i>	99
4.3.2 Is the Effect of Investor Sentiment on Momentum Profits Robust?	101
4.3.3 <i>Investor Sentiment, Momentum, and Market States</i>	101
4.3.4 <i>Investor Sentiment, Momentum, and Trading Volume</i>	105
4.3.5 <i>Is it a Size Effect?</i>	107
4.3.6 <i>Is it Risk?</i>	108
4.3.7 <i>Is the CB Sentiment Index Forecasting Future Macroeconomic Conditions?</i>	110
4.3.8 <i>An Alternative Sentiment Index</i>	111
4.3.9 <i>Alternative Lags for Optimistic and Pessimistic Sentiment</i>	113
4.4 Momentum Profits, Investor Sentiment and Long-Run Returns.....	113
4.5 Further discussion and concluding remarks.	115
5. Decision heuristics and Bayesian updating	138
5.1 Introduction	138
5. 2 Experimental Design	143
5.2.2 <i>Characterizing Attitudes Towards Risk</i>	144
5.2.3 <i>Eliciting Subjective Beliefs</i>	144
5.2.4 <i>Estimating the exact effect of strength and weight on elicited beliefs</i>	148
5.3 Results	151
5. 3.1 <i>Elicited Beliefs From The Raw Data</i>	151
5. 3.2 <i>Joint Estimates of Beliefs and Preferences Towards Risk</i>	155
5. 3.3 <i>Eliciting beliefs after controlling for preference towards risk</i>	160
5.3.4 <i>Estimating maximum likelihood parameters of the effect of strength and weight on elicited beliefs</i>	163
5.4. Further Discussion.....	166
5.5. What do these results mean for financial markets?.....	167
5.5 Conclusion.....	169
Appendix 5.1: Instructions for the Risk Aversion task	183
Appendix 5.2: Instructions for the Belief Elicitation task.....	184
Appendix 5.3. The Demographics questionnaire completed by subjects.	189

6. Conclusion.....	192
6.1 Caveats	194
6.2 Implications for future research	197
6.3 Closing note.....	199
7. Appendices	200
7.1 Heterogeneous expectations & behavioural asset pricing.....	200
7.2 Do investors systematically over or under react?.....	207
7.3 Further literature on market anomalies	210
7.3.1 <i>Price Momentum</i>	210
7.3.2 <i>The B/M effect</i>	211
7.3.3 <i>Post-Earnings announcement drift</i>	212
7.3.4. <i>Individual trading</i>	214
References	218

List of Tables

Chapter 3: Ambiguity aversion and analyst forecasts

3.1 Descriptive statistics for analyst forecasts.....	80
3.2 Ambiguity test.....	82
3.3 Market Impact of Forecasts.....	85
3.4 Market Impact of Forecasts and Market P/E ratio.....	86
3.5 Multivariable regressions.....	87
3.6 Returns Around Earnings Announcements.....	88

Chapter 4: Investor Sentiment and Price Momentum

4.1 Descriptive Statistics for Investor Sentiment Index.....	118
4.2 Momentum Profits Conditional on Investor Sentiment.....	119
4.3 Momentum Profits Conditional on Investor Sentiment and Market States.....	122
4.4 Momentum Profits Conditional on Investor Sentiment orthogonal Market returns.....	124
4.5 Regression of momentum profits on market states and investor sentiment.....	127
4.5 Momentum Profits Conditional on Investor Sentiment and Trading Volume.....	128
4.6 Momentum Profits Conditional on Investor Sentiment and Firm Size.....	129
4.7 Risk adjusted Momentum profits Conditional on Investor Sentiment.....	130
4.8 Momentum Profits Conditional on Investor Sentiment orthogonal future macro and VIX.....	131
4.9 Momentum Profits Conditional on an alternative Investor Sentiment index.....	132
4.10 Momentum Profits Conditional on different specification of sentiment.....	133
4.11 Long run profits of momentum portfolios conditional on sentiment.....	134

Chapter 5: Decision Heuristics and Bayesian Updating

5.1 An example of a lottery pair.....	173
5.2 The table in which subjects place their bets.....	175
5.3 Descriptive statistics for dice patterns.....	176
5.4 Maximum Likelihood estimates of Risk Neutral Subjective Beliefs.....	178
5.5 Comparison of bets between a risk neutral and a risk averse agent with the same beliefs.....	179
5.6 Maximum Likelihood estimates of risk aversion.....	180
5.7 Maximum Likelihood estimates of Risk Averse Subjective Beliefs.....	181
5.8 Maximum Likelihood estimates of coefficients of strength, weight and risk aversion.....	182

Acknowledgments

I greatly acknowledge the contribution of my supervisors Professors Daniel Read and Morten Lau. Their guidance, knowledge and support made this thesis possible.

I thank Professors Daniel Read and Glenn Harrison for using their funding resources to conduct the decision making experiment in Chapter 5.

I sincerely acknowledge the contribution of Professors Glenn Harrison, John Doukas and Avaniidhar Subrahmanyam. Being leading scholars in their fields they have immensely improved the quality of the thesis, and I am forever grateful for their involvement.

I want to thank my family for their endless love and support. I could not have done it without them.

Finally, I want to dedicate this thesis to Aina Brunet Moreno. During these years she provided balance by tirelessly reminding me of what really matters. This thesis is as much mine as it is hers.

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

Summary of contributions

Neoclassical finance applies rational expectations theory in the pricing of risky assets. However, in the past three decades the predictions of neoclassical theories have been challenged, leading to the emergence of “alternative” schools of thought. Behavioural Finance is one such school, which relaxes the assumption that investors consistently use Bayes law when updating their expectations. It draws motivation from psychology and behavioural economics, in an attempt to identify the behavioural determinants of portfolio choice.

This thesis builds on this theme and examines deviations from the Bayesian model that arise due to *i*) ambiguity aversion, *ii*) investor sentiment and *iii*) decision heuristics. These tests contribute to the literature by highlighting behavioural factors that affect portfolio choice that are omitted by neoclassical theories.

The motivation of the third chapter stems from the fact that, contrary to the requirement of Bayesian Updating, investors do not have *all* the necessary information at their disposal in order to estimate a single generating process for expected returns. Rather, investors are endowed with *noisy* information signals that relate to this unknown distribution in a *latent* way, and it is likely that in certain conditions they will not be able to pin down a *single* probability distribution. Such conditions where different probability distributions are possible are referred to as ambiguity, and a large literature in behavioural economics shows that in such circumstances peoples’ behaviour violates Bayesian Updating (see Camerer and Weber 1992 for a review). This finding has been confirmed in several experimental studies and has been used in theoretical models to explain stock market anomalies, such as the equity premium and excess volatility (e.g., Epstein and Schneider 2008). However, no study examines whether the behaviour of aggregate stock prices exhibits evidence of ambiguity

aversion. The third chapter attempts to fill this gap, by examining whether the pricing of analyst earnings forecasts shows traces of ambiguity aversion.

The fourth chapter relates to a relatively new stream of literature in Finance that examines the effects of investors' mood and feelings on their trading decisions. A behavioural literature shows that people in positive sentiment make optimistic choices, relative to objective probabilities, and people in negative mood make pessimistic ones. Because sentiment provides a measure of the propensity of investors to form erroneous expectations numerous studies examine whether it relates to stock market anomalies in an attempt to identify whether they relate to investors behavioural biases. The fourth chapter extends this literature and examines the relationship between investor sentiment and price momentum, an anomaly that has received both rational and behavioural explanations.

The analysis in the fifth chapter draws on psychological studies which show that deviations from Bayesian Updating arise because people use heuristics that are triggered by the characteristics of the information used to form subjective beliefs. This chapter examines the validity of one such important hypothesis proposed by Griffin and Tversky (1992) using rigorous experimental economics techniques. This chapter, apart from making several methodological advances, contributes to the literature by back testing the validity of a behavioural theory that has been extensively used in asset pricing theories (e.g., Barberis, Shleifer and Vishny 1998), shedding further light on the process by which information is transferred into asset prices.

The results from all three empirical chapters suggest that deviations from the Bayesian model can arise due to ambiguity aversion, investor sentiment and decision heuristics. Therefore, the thesis expands the behavioural literature by highlighting further behavioural factors that affect asset prices, conveying the message that the neoclassical definition of risk is too simplistic to consistently capture the behaviour of asset prices.

The following section provides a more detailed discussion of the relevant literature, whilst developing the hypothesis tested in each empirical chapter.

1. Introduction

1.1 A shift of paradigm in financial economics

The Neoclassical paradigm states that asset prices fully reflect asset risk. This prediction emerges in an environment where investors analyze correctly all the available information that concerns financial assets' risk and return profile, and take positions according to their risk tolerance. This process determines the equilibrium price of the asset. Changes in price occur *only* when new information arrives in the market, which alters the asset's profile. This alteration prompts agents to rebalance their portfolios, leading to a new equilibrium price. Therefore, since asset prices always incorporate correctly all available information about asset risk and return, price changes are unpredictable from available information; hence markets are informationally efficient [Samuelson (1965), Fama (1965)].

This paradigm evolved during the 20th century, a period where economic models were based on solid microeconomic foundations, and derived their conclusions with mathematical rigor. However, rigorous mathematics came at a cost. They quickly became intractable once the complexity of financial markets was recognized. Thus, various simplifying assumptions were necessary to allow manageable calculations and closed forms solutions. These were that investors are fully rational. They have the resources and computing capacity to analyze all the available information correctly and infer assets' true risk and return profile. In addition, asset markets are frictionless, without trading costs or taxes, and the information required to estimate the distribution of asset returns is costless and simultaneously available to all market participants.

Clearly these assumptions are unrealistic. But, as advocated by Friedman (1952) in his famous essay, a theoretical economic model *should not* be assessed based on the realism of its assumptions, but on whether it can predict with a reasonable degree of accuracy the

quantity in question. If it can, then the assumptions should not come under scrutiny because they are unrealistic.

So, can asset pricing theories and the Efficient Market Hypothesis explain the behaviour of asset prices? It seems that they cannot, at least not conclusively. A large literature has accumulated over the past three decades that highlights patterns in the data that contradict the notion that prices reflect systematic risks and are unpredictable. This literature on “market anomalies” can be broadly classified into three categories that relate to: *i*) the time series behaviour of prices, *ii*) the cross sectional behaviour of prices, and *iii*) the general characteristics of financial markets.

In terms of the time series behaviour of prices Fama and French (1987) and Lo and McKinlay (1988), amongst others, show that prices exhibit positive serial correlation. Given that the information flow does not exhibit serial correlation this finding is difficult to explain in a rational and completely frictionless market.

More powerful evidence against the neoclassical model arises from the cross sectional behaviour of prices, where various factors that cannot uncontroversially be classified as sources of systematic risk predict returns. For example, Lakonishok, Shleifer and Vishny (1994) and Fama and French (1992) show that returns relate to book value to market value ratios, and Banz (1981) and Fama and French (1992) to firm size. Daniel and Titman (1997) in a very important study show that these patterns are inconsistent with covariance-driven risk-premia. In addition, prices are predictable after various events, such as earnings announcements [Livnat and Mendenhall (2006)] and analyst earnings forecasts [Givoly and Lakonishok (1979), Gleason and Lee (2003)]. Simple strategies that exploit these phenomena consistently earn high abnormal returns, which are unrelated to asset risk as implied by the CAPM. Further, future returns are reliably predictable *solely* from past returns. DeBondt and Thaler (1985), and a large literature thereafter, document long run *reversals*. That is, over

long horizons of two to five years, past losers consistently outperform past winners. Thus, a strategy that is long on losers and short on winners, earns significant abnormal returns relative to the CAPM. Jegadeesh and Titman (1993) demonstrate a medium horizon *momentum* effect. That is, over periods of three to twelve months past winners outperform past losers, making a strategy that is long on the former and short on the latter profitable.¹

This evidence from the cross sectional behaviour of returns contradicts the notion that prices cannot be predicted from publicly available information.

Lastly, some of the overall characteristics of the market seem inconsistent with the neoclassical model. The volatility of fundamentals is not sufficient to explain the volatility of prices, [Shiller (1981)], and daily trading volume is far too great to be solely driven by risk-based rebalancing decisions [Odean (1999)]. In addition, historically stock returns have displayed a premium that cannot be easily reconciled with the amount of covariance risk they entail [Mehra and Prescott (1985)].

Because of these inconsistencies the paradigm that governs the practice of Financial Economics in the last two decades has experienced a major shift. Economists, dissatisfied with the inability of neoclassical theories to explain these phenomena, endeavour to create asset pricing models that relax the more demanding neoclassical assumptions, in an attempt to better understand how financial markets *actually* behave. This will lead to more accurate predictions, and ultimately increase the welfare of market participants.

1.2 Behavioural Finance

Behavioural Finance is an alternative school of thought that relaxes the assumption that investors are fully rational utility maximizers. It draws from findings in experimental economics and psychology on how people behave in conditions of uncertainty, and develops

¹ See appendix 7.3 for further discussion on these anomalies.

models and hypotheses that examine whether the “anomalous” phenomena in financial markets are at least partly related to investors’ *bounded* capacity to process information and estimate risk and return.

This is the underlying motivation of this thesis. My objective is to draw on the literature in decision making and propose behavioural hypotheses that: *i*) highlight conditions in which investors behaviour can depart from the assumption of perfect rationality, and *ii*) examine the effect of these departures on asset prices. These tests expand the field of behavioural finance by increasing our understanding of the forces that operate in financial markets, and ultimately the price formation process.

In order to describe the general framework in which these hypotheses are developed it is useful to define what exactly is meant by the term “rational” behaviour. Rationality relates to the way investors form expectations. If, on average, investors’ expectations accurately describe the future performance of the economy (or a company), and if these expectations are used to make choices that maximize their utility, their choices are rational.

The statistical model, which provides the basis for rational expectations, is Bayes rule. This is a method of combining prior knowledge with new information to form an updated expectation for the relevant quantity, i.e., asset prices. This rule is central in neoclassical theories, which view price changes as Bayesian responses to new information. Specifically, the assumption made by such theories is that investors correctly update their old expectations of the assets’ risk and return profile using new information, and through their trades set new, fully revealing equilibrium prices.

The decision making literature, however, has identified circumstances in which the behaviour of investors can *systematically* contradict Bayesian Updating. That is, in certain conditions investors may overestimate the likelihood of certain events and underestimate the likelihood of others, which leads to erroneous expectations. Since prices will reflect *these*

erroneous expectations they will not be set relative to the fundamentals and thus will temporarily be out of balance. With the passage of time fundamentals become transparent and reveal the mispricing; therefore prices gravitate to their correct levels. This process of correction, however, amounts to price predictability that is incompatible with the Efficient Market Hypothesis.

1.3 Hypotheses and contributions

As mentioned earlier the underlying motivation of the thesis is to draw on the literature in decision making and propose hypotheses that examine whether in certain cases the behaviour of investors deviates from the optimal Bayesian model. This section outlines the hypotheses examined, and their contribution to the literature.

1.3.1 Ambiguity in analyst forecast accuracy

Neoclassical theories assume that investors can gather all the available information and estimate a *single* joint (Bayesian) distribution of future asset returns. Because the relationship between information and expected returns is so complicated, it is difficult in practice to always infer the *latent* distribution of returns. In some cases different probability distributions can be possible, and investors will be unable to confidently distinguish between them (as explained in Epstein and Schneider 2008 among others). The condition that many different distributions of future returns can be possible is called ambiguity, and a large experimental and theoretical literature shows that agents are particularly averse to ambiguity [see Camerer and Weber (1992) for a review of the evidence]. Specifically, when different distributions are possible investors form *pessimistic* expectations, behaving *as if* the *worst-case* distribution that may arise is the most likely to occur. This is inconsistent with the

Bayesian model, which predicts that agents have at their disposal the complete information set and are always able to infer the correct empirical distribution.

In the third chapter we examine whether ambiguity aversion affects the way investors update their earnings expectations using analyst forecasts, when the accuracy of these forecasts is ambiguous. Neoclassical theories suggest that investors are always able to ex-ante infer the distribution of forecast accuracy for each forecast, and thus respond more strongly to forecasts that are expected to be accurate [Abarbanell, Lanen and Verecchia (1995)]. This requires that investors fully comprehend the variation in forecast accuracy, an arguably difficult task. If investors, in certain cases, cannot confidently determine an exact distribution of forecast accuracy, the concept of ambiguity can be relevant. Based on the literature in behavioural economics we propose the *pessimism-and-correction* hypothesis, which states that insofar the accuracy of the forecast is ambiguous, investors will respond pessimistically and set prices too low. As the future unfolds the mispricing becomes apparent, prices will rise to exhibit an ambiguity premium.

This chapter contributes to the literature in the following way: Whilst previous examinations of ambiguity have either been experimental or theoretical, this is the first study that provides a definition of ambiguity that can be taken into the “field” and examines whether ambiguity aversion operates in the real-world of financial markets. This is important, because ultimately the applicability of behavioural theories is validated with such empirical tests.

In addition, this study extends the literature that analyzes investors responses to analyst forecasts. Traditionally, the responses of investors to analysts forecasts has been assumed to be in accordance to Bayes law [Abarbanell, Lanen and Verecchia (1995)], but recently studies have shown that behavioural biases affect the way investors price analysts forecasts [Sorescu and Subrahmanyam (2006)]. This study proposes a new behavioural

hypothesis, which aims to examine whether ambiguity and ambiguity aversion affect investors responses to analyst forecasts.

The results provide support to the pessimism- and-correction hypothesis. Conducting a two-stage event study we find that when the accuracy of the forecast is ambiguous, the initial response is to update the expectations pessimistically and set prices too low. Then, during an adjustment period, prices rise to correct this initially pessimistic response.

1.3.2 The effect of Investors' Sentiment on Price Momentum

Another source of bias that can make investors' decisions incompatible with Bayes rule is their mood or sentiment at the time of the decision. A large literature in decision science finds that when people in a good mood make excessively optimistic decisions, whereas people in bad mood make excessively pessimistic ones [Arkes, Herren and Isen (1988), Bower (1981), (1991); Wright and Bower (1992), among others].

This finding is particularly relevant for financial theories, because prices reflect investors' expectations. If investors are prone to sentiment, prices will exhibit their excessively optimistic and pessimistic views, and thus will become predictable as sentiment dissipates and prices slowly correct. Indeed using various proxies for investors' sentiment, the literature has produced evidence that following periods of high sentiment asset returns are lower [Hirshleifer and Shumway (2003), Edmans, Garcia and Norli (2007), Lemmon and Portniaquina (2006), Baker and Wurgler (2006); (2007) among others].

The evidence that sentiment is priced suggests that investors' mood affects the pricing of information in financial markets. In the fourth chapter we extend the literature by examining whether this finding affects the performance of the momentum trading strategy, a phenomenon firstly documented by Jegadeesh and Titman (1993), which has received both rational and behavioural explanations.

Since sentiment has been shown to affect perceptions of future contingencies, it can affect the performance of the momentum strategy. The fourth chapter conditions momentum profits on sentiment, thus sheds light on whether this market anomaly has behavioural foundations. Particularly, sentiment underlies the behavioural theory of momentum proposed by Daniel et al (1998), which states that because investors' expectations are miscalibrated, their trading actions generate momentum. It follows that when investors are optimistic, their beliefs will be even more miscalibrated relative to Bayesian probabilities, therefore the predictions of Daniel et al (1998) will be magnified. Specifically, the *optimism hypothesis* predicts that short run momentum will be higher when investors are optimistic, and these momentum portfolios, since they reflect overly optimistic expectations, will experience reversals in the long run.

Examining this hypothesis contributes to the literature in the following way: Price momentum is the only CAPM anomaly not explained by the Fama and French (1993) three factor model, and a major source of controversy in the literature. Some believe that the phenomenon reflects investors' behavioural biases [Daniel, Hirshleifer and Subrahmanyam (1998), among others], whilst others that it reflects compensation for bearing systematic risk [Chordia and Shivakumar (2002)]. This analysis in this chapter provides evidence that can relieve some of this tension, by examining whether behavioural theories fit the data better than rational theories.

Secondly, it adds to an expanding literature that examines the effects of investor sentiment in financial markets. Thus far, investor sentiment has been linked to the post earnings announcement drift [Livnat and Petrovic (2008)], corporate disclosure [Bergman and Roychowdhury (2008)], IPO's [Cornelli, Goldreich and Ljungqvist (2006)] and the size effect [Baker and Wurgler (2006; 2007)]. These studies are important because they help explain various anomalous phenomena in the market. This study extends this literature by

analyzing the relationship between investor sentiment and momentum, another puzzling stock market anomaly.

Using the consumer confidence survey compiled by Conference Board to proxy investor sentiment, the results provide strong support for the optimism hypothesis. We find that momentum is *only* significant when investors are optimistic, and that these portfolios exhibit strong long run reversals. This result is robust to different specifications of investor sentiment, an alternative sentiment index, controls for company size and trading volume, market states, microstructure biases and risk adjustments, and is consistent with the view that momentum and reversals arise jointly from investors' optimistic biases.

1.3.3 Decision Heuristics

A large literature, principally pioneered by psychologists Daniel Kahneman and Amos Tversky, demonstrates that people in certain cases tend to make decisions using heuristics that lead to systematic deviations from the Bayesian model. For example, Edwards (1968) finds that in the presence of new information people cling excessively on their priors, whereas Tversky and Kahneman (1971) find the reverse. In a very influential study Griffin and Tversky (1992) suggest that these biases commonly arise as functions of the characteristics of the information used by decision makers to update their expectations. Specifically, when the information is salient but low in credence (high strength low weight), people *overreact* as suggested by Tversky and Kahneman. When the information is moderate and high in credence (low strength high weight) people *underreact* as suggested by Edwards.

The study by Griffin and Tversky (1992) can explain the evidence that in some cases prices overreact to information [De Bondt and Thaler (1985)], whereas in others they underreact to information [Jegadeesh and Titman (1993)]. This addresses a very important criticism voiced by Fama (1998), that in order for behavioural theories to be successful they

must be able to predict when investors systematically over or under react. Indeed one of the most important behavioural asset pricing theories proposed by Barberis, Shleifer and Vishny (1998) explains the sort-run momentum and long run reversal effects by appealing to the hypothesis proposed by Griffin and Tversky (1992).²

One limitation of these studies, however, is that the robustness of the Griffin and Tversky (1992) results has not been established. This is particularly worrying given the evidence that the behavioural biases documented by psychologists can *disappear* once the experimental design carefully addresses all the necessary requirements of a controlled experiment as outlined by Smith (1982) [e.g., as in Grether (1980)].

In the fifth chapter we address this issue and examine the robustness of the *strength-weight hypothesis* proposed by Griffin and Tversky. We design a novel experiment in which subjects are properly incentivized to make correct decisions. Their beliefs are extracted from a specific betting ‘game’ that avoids the use of the terms probabilities as in Fiore et al (2008). Further, decisions are based on real, observable information, and not hypothetical events. Lastly, we allow subjects to be risk averse, as opposed to assuming that they are risk neutral as is common in the literature, and specify a maximum likelihood model whereby all the parameters are estimated jointly (as in Andersen et al 2008).

This chapter contributes on the literature in the following way: Firstly, the experiment conducted offers out of sample evidence on whether the Griffin Tversky hypothesis captures decision making better than Bayesian updating does, and thus can *justifiably* be used to explain the apparent over and under reactions in prices. In addition, we make several methodological contributions to the literature. We design a novel experiment that allows extracting subjective beliefs and testing Bayesian updating. In addition, we allow subjects to be risk averse in a framework that all parameters are estimated jointly.

² Sorescu and Subrahmanyam (2006) also test an over-under reaction hypothesis based on the results of Griffin and Tversky (1992).

The results broadly support the predictions of the strength-weight hypothesis, thus justify the use of this theory in explaining over and under reactions in financial markets. However, we find that the magnitude of the bias is significantly *reduced* in our study, compared to the results of Griffin and Tversky, especially when we allow for risk aversion. This finding suggests that it is prudent to establish the robustness of behavioural biases prior to applying them in economic models.

All three empirical chapters demonstrate that deviations from Bayesian updating occur systematically amongst investors, due to ambiguity aversion, investor sentiment and decision heuristics. Since these behavioural characteristics are omitted from neoclassical theories that assume Bayesian updating, the evidence provided in the thesis highlights that “rational” theories are too simplistic to consistently capture the process by which expectations are updated and information is transferred into asset prices. Rather, this evidence supports the notion that the “anomalies” we observe in the marketplace, defined vis-à-vis the rational theories, are at least partly related to investors’ inability to consistently update their expectations using Bayes rule.

Chapter 2 in the thesis explains in more detail the mechanics of the Bayesian model, and the discrepancies that can arise from investors’ behavioural characteristics. Chapters 3, 4 and 5 present the empirical chapters, chapter 6 concludes the thesis and chapter 7 presents several appendices.

2. Bayes Rule and Investment decisions

This chapter explains the mechanics of asset pricing based on Bayes Rule, outlines what these mechanics imply for real financial markets, and details possible behavioural deviations.

2.1 Bayesian asset pricing

Each time period new information signals are added to the information set. The quality of the expectations that are formed using this ever-expanding information set depends on the extent that the market has comprehended the implications of the new ‘data’ on the future streams of cash flows.

In the neoclassical paradigm agents can correctly analyze the information set, and apply Bayes rule to update their expectations. Thus asset prices are unpredictable. The following section details the mechanics of the Bayesian model.³

Assume that:

- One asset that pays 100% of its earnings as dividends at some terminal date.
- Assume one representative, risk-neutral agent whose earnings expectations determine the market price as $P = F(E^e)$.

Suppose that at time $t-1$ the agent, after observing the available information believes that the earnings generating process $G(E)$ is as follows: $G(E) \sim N(\bar{E}, \frac{1}{\rho_{\bar{E}}})$. Thus the agent

believes that earnings are (normally distributed) with a mean value \bar{E} and variance

$Var(\bar{E}) = \frac{1}{\rho_{\bar{E}}}$. The parameter $\rho_{\bar{E}}$ in this case is the Bayesian notation for precision. This will

³ These equations demonstrate standard updating for normally distributed variables. See Anderson (1984), Chap. 2 for an exposition.

be explained in more detail shortly. Having these expectations, $E(E) = \bar{E}$, the agent sets the asset price $P(\bar{E}) = P_0$.⁴

Then suppose that at time t the agent observes a noisy information signal $S(E)$ that suggests a different distribution of earnings, as shown below:

$$S = E' + u \quad u \sim N(0, \sigma_u^2) \quad E' \sim N(\tilde{E}, \sigma_{E'}^2) \quad (2.1)$$

It is useful to think of this signal as the product of the investor's analysis of the available information. For example, the investor has gathered some information about company earnings, and has produced the signal $S(E)$ that suggests a different distribution for the latent variable E . The signal is noisy because the information gathered is not exact. It contains a substantive component that relates to future earnings, but also a noisy component that is unrelated to fundamentals. From the normality of E' and u it follows that the expectation of earnings based on this new signal is $E(S) = \tilde{E}$ and the precision $\frac{1}{P_s} = \sigma_{E'}^2 + \sigma_u^2$.

Assume that $\tilde{E} > \bar{E}$. That is, the information signal is “good news”, and suggests that earnings will be higher than expected. Bayesian updating requires that the agent optimally combines his priors, $G(E)$, and the new information, $S(E)$ to arrive at a new expectation for earnings, $E(E/S)$. Using Baye's theorem the mean of the new conditional distribution of earnings is equal to:

$$E(E/S) = \bar{E} \left(\frac{P_{\bar{E}}}{P_{\bar{E}} + P_s} \right) + \tilde{E} \left(\frac{P_s}{P_s + P_{\bar{E}}} \right) \quad (2.2)$$

This new conditional expectation will reflect that price at time t , P_t . This price is a precision weighted average of the prior expectation, \bar{E} , and the expectation implied by the new signal, \tilde{E} . Rearranging and subtracting $E(E)$ yields:

⁴ The price is set only according to the mean of the distribution $G(E)$ because the agent is risk neutral.

$$E(E/S) - E(E) = (\tilde{E} - \bar{E}) \left(\frac{p_s}{p_s + p_{\bar{E}}} \right) = P_1 - P_0 \quad (2.3)$$

$E(E/S)$ is the new expectation of the agent given the observance of the signal and the previous unconditional expectation. Their difference between $E(E/S)$ and $E(E)$ reflects the magnitude of the update to the expectations induced by the signal, and thus the change in prices, $P_1 - P_0$. From the above formula it is shown that according to Bayes rule, the extent the new signal affects expectations and asset prices increases with: *i*) the amount of new information that it brings to the market, $\tilde{E} - \bar{E}$, and *ii*) its precision, ρ_s .⁵ Given that the agent in the Bayesian model is fully rational these dimensions are inferred correctly from the available information, and the new price P_1 is determined (which is unpredictable at time $t+1$).

The fact that the agent is assumed to be risk neutral does not affect the way information is used in the Bayesian model, which is what this analysis aims to highlight. If risk attitude is accounted for the price of the asset at any time period will reflect the expected utility of the expectation, i.e., $P_0 = U[E(E)]$ and, insofar agents are risk averse, will carry a risk premium.

2.2 Bayesian asset pricing in practice

The information in neoclassical Bayesian models is essentially an objective statistical distribution. In practice, however, investors only observe information signals that relate to this distribution in a latent manner. For example, investors observe information in the form of newsletters, corporate announcements, analyst forecasts, macroeconomic indicators, rumours, word of mouth etc. From these signals investors' attempt to estimate the assets risk and return

⁵ The update is also affected by the precision of the priors p_E , but for the analysis only focus on the effect of the new information.

profile. Importantly, these signals vary in their informativeness about this (latent) profile, as shown by Equation 2.1. That is, whilst they contain a substantive component that is related to fundamentals, they also contain a noisy component unrelated to fundamentals. The neoclassical assumption, therefore, implies that investors observe all the relevant information signals, extract their substantive component and estimate the Bayesian distribution.

This task is extremely complicated. Black (1986) in a seminal paper was the first to point out that investors' will not be able to perform this task successfully because disentangling substance from noise can be daunting. Therefore, in some cases, investors will trade on information that *they* believe is substance, but in fact it will be noise. Furthermore, the evidence from psychology and behavioural economics suggests that in such complicated environments such as the stock market, where the latent processes are extremely difficult to comprehend and feedback is slow and inconclusive, the effect of behavioural biases on the disentanglement of substance and noise becomes more important [Hirshleifer (2001)].

The next section categorises and describes the main behavioural biases that have drawn the attention of financial economists.

2.3 Behavioural asset pricing

Hirshleifer (2001) groups the behavioural biases that may operate in financial markets into three categories: Biases arising from: *i*) heuristic simplification, *ii*) self-deception and *iii*) biases arising from investors' sentiment. Since in the thesis we refer to a bias as a deviation from the neoclassical model we add a fourth category as, *iiii*) biases that arise from investors' preferences that do not conform to the axioms of Expected Utility.

A complete discussion of all the behavioural biases and how they contribute to non-Bayesian updating is beyond the scope of the thesis. Rather, the focus will be on the behavioural biases that are empirically examined, by providing a literature review that

explains the main findings and highlights the gaps that are tackled by the thesis in each chapter.

2.3.1 Heuristic simplification -Representativeness

A robust phenomenon firstly documented by Amos Tversky and Daniel Kahneman (1974) is Representativeness (henceforth R). According to this decision rule people assess the likelihood of a hypothesis by the extent that the available information retains certain salient characteristics that are exhibited by the cases where the hypothesis holds. A classic experiment conducted by Kahneman and Tversky that illustrates R is the following: They provided their subjects with information about a hypothetical person named Linda: *“Linda is 31 years old, single outspoken and very bright. She majored in philosophy. As a student she was deeply concerned with the issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.”* Then subjects were asked to rank possible alternatives about Linda according to their probability, including event A that *“Linda is a bank teller”* and event B that *“Linda is a bank teller and is active in the feminist movement”*. Obviously the latter event is a subsumed by the former, so its probability cannot be greater. However, Kahneman and Tversky (1972, 1973) find that 85% of the subjects ranked event B as more probable than A. It seems that because the description offered for Linda is *representative* of a feminist, makes subjects insensitive to the statistical fact that the frequency of bank tellers is larger than the frequency of bank tellers that are feminists.

The effects of R in the model described above can be captured if the investor makes his expectations based on the signal, \tilde{E} , more extreme by a factor $\lambda > 1$. Thus, whilst the correct expectation is $E(E) = \bar{E}$, the agent believes that the new signal implies $\lambda \tilde{E}$. If the signal is good news, and $\tilde{E} > \bar{E}$, the price update will be affected by λ as $\lambda P_I > P_I$, and will

thus be more extreme. Because the investor, when updating his expectation, has overweighted the signal, the price of the asset λP_t will be too high relative to the fundamental value (by the proportion λ). As the future unfolds the mispricing becomes apparent, and prices revert to fundamental values. This however induces price predictability since the behavioural bias of representativeness has induced serial correlation in prices.

Representativeness has been incorporated into asset pricing theories. For example, in Barberis et al (1998), an investor observes a sequence of good (or bad) news for a company. This makes the investor who displays R believe that this good performance is likely to continue in the future, failing to recognize that most companies cannot grow (or lose) indefinitely. This makes the stock of the company mispriced, inducing a period of correction, as the price slowly returns to fundamental value. R is also studied by Rabin and Vayanos (2009) in a theoretical asset pricing context, and is able to explain a number of anomalies including fund-flow puzzles and the presence of momentum and reversal in returns.

The empirical literature has produced evidence that support the existence of representativeness. DeBondt and Thaler (1985, 1987) show that stocks that have been performing poorly during the last 3-5 years, consistently outperform companies that have been performing well during the same period. Lakonishok, Shleifer and Vishny (1994) show that companies that have experienced extreme sales growth underperform those that have experienced poor sales growth. LaPorta (1996) shows that companies that are associated with high forecasts for long run growth issued by security analysts, also experience lower average returns. This evidence suggests that the market becomes very excited when it analyzes *representative* information for companies that are performing extraordinarily well or extraordinarily poor, and omits to recognize that the performance of the company is likely to regress to “normal” performance. Therefore, the market exhibits the fallacy of representativeness, and extrapolates these trends into the future, overvaluing the winners and

undervaluing the losers. With the passage of time, as these mispricings becomes apparent, prices drift to fundamental values, with past “losers” outperforming past “winners”.

A caveat of these studies however, is that the robustness of decision heuristics such as representativeness has not been fully established in the literature. If these heuristics are not genuine their application in financial markets is erroneous. To illustrate, representativeness has been identified in psychological experiments conducted by Kahneman and Tversky (1971, 1972). Economists, however, generally oppose the methods used by psychologists such as Kahneman and Tversky. Grether (1980, 1992) addresses this debate and tests whether R is found in experiments that satisfy the necessary experimental conditions (these conditions are explained in more detail in chapter 5). His results show that the evidence of R is much weaker, casting doubt on the robustness of the relation between decision heuristics identified by psychologists and asset pricing anomalies. In the fifth empirical chapter the thesis addresses this issue and, similarly to Grether, back-tests an important behavioural theory proposed by Griffin and Tversky (1992), which has been used to explain the behaviour of asset prices, both theoretically (Barberis et al 1998) and empirically (Sorescu and Subrahmanyam 2006).

2.3.2 Self deception- Overconfidence

Overconfidence is the tendency to place an irrationally excessive degree of confidence in ones abilities and beliefs, and is manifested in two ways; the better than average effect, and miscalibration.⁶ According to the former investors are overconfident because they believe that they are better at selecting stocks relative to their peers. According to the latter they are overconfident because the confidence in their beliefs is higher than what

⁶ Overconfidence can also arise from heuristic simplification, as suggested by Griffin and Tversky (1992). In fact overconfidence is the end product of any non-Bayesian decision rule since the decision maker holds erroneous beliefs which are believed to be correct.

is merited by objective probabilities. Both types of overconfidence imply that investors *falsely* believe that the investment signals they have generated are of better quality than what they really are. Classic studies that demonstrate that individuals are overconfident include Oskamp (1965), Alpert and Raiffa and Lichtenstein, Fischhoff and Philips in Kahneman et al (Eds) (1982).

In the Bayesian model outlined above an overconfident investor constructs a signal that he believes is more accurate than what it really is, inflating its precision by an overconfidence factor $k > 1$. This means that whilst the true precision of the signal is p_s , the investor who displays the fallacy of overconfidence believes that the precision of the signal is greater than what it really is, i.e., $kp_s > p_{ss}$. Since, as shown by Equation 2.3, the impact of the signal on expectations depends on its (perceived) accuracy, an overconfident investor will overweight his private signal, as he believes that its accuracy is better than what it really is.

Overconfidence has received considerable attention in the finance literature, both theoretically and empirically. Odean (1998) in a theoretical study, demonstrates that overconfidence can lead to serial correlation in prices, and explains the high trading volume and volatility in stock markets. Same conclusions are drawn in the theory proposed by Daniel, Hirshleifer and Subrahmanyam (1998).

These predictions have received support from the empirical studies. Odean (1999) analyzes the trading behaviour of a large number of individual investors and comes to the conclusion that they trade excessively and, consequently, lower their returns. Similar results are reported by Barber and Odean (2000) who conclude that investors' positions do not generate enough return to cover the transaction costs. In a rational framework whereby an investor correctly anticipates the future return of the asset, a trade takes place insofar the expected benefit offsets the cost. Odean and Barber suggest that their findings are consistent

with overconfidence, whereby investors believe that their signals are more accurate than what they really are, and thus overweight them when they update their expectations and trade.

Barber and Odean (2001) further test the prediction that overconfident investors will trade excessively by examining the trade behaviour of males and females. The prediction is that, since males are in general more overconfident than females [see references in Barber and Odean (2001)], they will trade more and realize lower returns. Their results confirm this hypothesis. The finding that overconfidence investors trade more heavily is also supported by Glaser and Weber (2009).

A noteworthy study in the field is conducted by Grinblatt and Keloharju (2008). They have access to the entire database of trades in the Finnish market, as well as information on the characteristics of traders, such as their aptitude, sex etc. This information about investors' aptitude comes from various tests and questions that all males in Finland answer when they enrol for their compulsory military service. Grinblatt and Keloharju measure overconfidence by regressing the aptitude of investors (based on their IQ test) on their perceived self-reported ability and use the residuals from this regression as their overconfidence measure. This measure is a cleaner proxy for overconfidence than the measures used in the previous studies, such as gender. The results of Grinblatt and Keloharju also support the overconfidence hypothesis, as they show that overconfident investors trade more aggressively.

The application of overconfidence in financial markets has provided a behavioural explanation for the excessive volatility and high trading volume that is observed. Furthermore, models of overconfidence can also explain serial correlation in prices.⁷

⁷ With no other assumptions overconfidence leads to negative serial correlation. However, with additional assumptions it can also explain positive serial correlation [Daniel et al (1998), Odean (1998)]. For example, in the static version of the Daniel et al (1998) model the effect of the private signal on asset prices is excessive due to overconfidence. When the terminal value of the asset is revealed prices return to fundamental values. If the private signal was positive (negative) it pushed prices too high (low), inducing an opposite adjustment in the next period, therefore negative serial correlation. However, when dynamic overconfidence is introduced due to self attribution, confirming news in a second round can push the prices further from fundamental values before returning to equilibrium, inducing first positive and then negative serial correlation.

2.3.3 Mood and Feelings- Investor sentiment

Sentiment refers to whether an individual, for whatever *extraneous* reason, feels excessively optimistic or pessimistic about a situation. A large body of the psychology literature finds that peoples' *current* sentiment affects their judgment of future events. For example, Johnson and Tversky (1983) show that people that read sad newspaper articles subsequently view various causes of death, such as disease etc., as more likely than people who read pleasant newspaper articles. In general, the evidence from experimental psychology shows that people with positive sentiment make optimistic judgments and choices, whereas people with negative sentiment make pessimistic ones [Bower (1981, 1991); Arkes, Herren, and Isen (1988); Wright and Bower (1992); among others].

This finding has been applied in financial markets. The general result is that after periods of high sentiment, stock returns are *lower*, especially for stocks that entail more subjective valuations. This is in line with the evidence from psychology that people with optimistic sentiment excessively inflate the likelihood of good future contingencies, and *vice-versa*. For example, Hirshleifer and Shumway (2003) use the psychological evidence that mood is related to the weather, and that particularly people are in a better mood during sunnier days [see references in Hirshleifer and Shumway (2003)]. They collect data on daily cloudiness and daily nominal return of stock market indices. Their results uncover a significant "weather effect", which supports the predictions of the sentiment hypothesis. For example, the annualized New York market return is a significant 24.8% on sunnier days, versus 8.7% on cloudy days. This finding is robust in 25 out of 26 cities used in the analysis, and thus it is unlikely to have arisen by chance.

In another interesting study Edmans, Garcia, and Norli (2007) capture investors' mood using soccer results, based on the findings that sports results have a significant effect on mood [see references in Edmans et al (2007)]. Using a cross section of 39 countries they

find that losses in international football games have an economically and statistically significant negative effect on the losing country's stock market. Particularly they show that the next day return of the stock market index is 38 basis points lower than the average, and that this decrease in market returns is unrelated to productivity or lost revenue issues. Further, the authors document this loss effect after international cricket, rugby and basketball games.

Kaplanski and Levy (2009) capture the effects of sentiment using aviation disasters. They argue that because such events are very salient and attract substantial and vivid media coverage, they cultivate negative sentiment amongst the investing public (and the general population), which they hypothesize will negatively impact on their future expectations. Consistent with this view Kaplanski and Levy find that the market records an average loss of 60 billion dollars per aviation disaster, which is fully recouped during the next trading days. This reversal indicates a temporary sentiment-related overreaction.

Although these studies do not by themselves resolve any significant puzzle they demonstrate one important result: That investors' sentiment affects their trading behaviour and feeds into prices in a manner that is *unrelated to fundamentals*.

Another stream of literature uses indices of consumer confidence to measure sentiment, such as the index compiled by the *University of Michigan*, the *Conference Board* and *Investor Intelligence*. The findings from such studies are in line with the aforementioned results. For example, Brown and Cliff (2005) use the *Investor Intelligence* index and find that this index is related to pricing errors derived independently from a valuation model, and negatively relationship to future returns. Similar results are shown derived by Lemmon and Portniaquina (2006).

Lastly, some studies extract sentiment proxies from market data that reveal investors' propensity to invest in the stock market. The intuition is that when investors engage in stock investing more aggressively, it is because they are feeling optimistic about the capacity of the

market to generate profits. Baker and Wurgler (2006) create a sentiment index from such sentiment revealing variables, such as trading volume, first day returns of IPO's the closed end fund discount. Frazzini and Lamont (2008) construct such a measure using flows of capital into mutual funds. Both studies demonstrate that returns are lower after periods of low sentiment.

These results suggest that investor sentiment does affect asset prices, contrary to the notion that investors form Bayesian beliefs and maximize their expected utility. Motivated from this finding several studies examine whether any of the market anomalies are due to sentiment. Baker and Wurgler (2006, 2007) explain the size effect using a sentiment variable constructed from market data. They find that after periods of pessimism smaller companies outperform larger ones, giving rise to the well known size premium identified by Banz (1981). However, this size premium *disappears* after periods of market optimism. Baker and Wurgler explain that sentiment affects the way investors respond to uncertainty. When they are pessimistic they tend to overweight the negative scenarios that may arise, therefore are repelled by smaller stocks that entail more subjective valuations. This reduces the prices of small stocks relative to larger ones, inducing a size premium in the next period. On the contrary, when investors are optimistic they are drawn to the uncertainty that surrounds smaller companies as they tend to overweight the optimistic scenarios that may arise, increasing their current price and lowering their future return.

In a recent study Livnat and Petrovic (2009) explain the post-earnings announcement drift anomaly firstly documented by Ball and Brown (1968) using sentiment (for details on this anomaly see appendix 7.3.5). These authors find that the drift is stronger for firms with positive surprises when sentiment is pessimistic. This is in line with the view that pessimists relatively disregard the good news in the earnings surprise because their future expectations are generally pessimistic; therefore they underweight to information content of the earnings

announcement. As the future unfolds and fundamentals are slowly revealed prices slowly gravitate toward their correct levels, which results to a more pronounced post-earnings announcement drift. On the contrary, when investors are optimistic the drift is less pronounced because the initial reaction to the good news is stronger, due to the fact that investors are expecting the economy to grow. These results suggest that the post-earnings announcement drift, a phenomenon coined by Fama as “the granddaddy of all anomalies”, is to an extent related to investors’ tendency to form erroneous expectations.

Cornelli et al (2006) highlight a link between investors’ sentiment and the long run underperformance of Initial Public Offerings (IPO’s), a phenomenon principally documented by Ritter (1991) and Loughran and Ritter (1995). Cornelli et al (2006) find that first day returns to IPO’s are positively related to an index of sentiment of the small investors. They also find that long run underperformance is much greater when the initial value of the index is high. This suggests that small investors become over-excited about new companies when they are optimistic, and overvalue their growth opportunities. As the future unfolds, however, these growth opportunities do not materialize therefore prices regress to fundamental values.

The application of sentiment in the context of the size effect, the post-earnings announcement drift and the long run underperformance of IPO’s has provided some support to the notion that these anomalies are at least partly driven by investors irrational behaviour. Since sentiment provides an indication of the extent to which investors’ expectations deviate from Bayesian, this approach of conditioning market anomalies on sentiment seems fruitful in terms of disentangling rational from behavioural explanations for various other anomalies.

Another market anomaly that has received both types of explanations is price momentum. Jegadeesh and Titman (1993) observed that stocks that performed well in the recent past (winners) continue to outperform stocks that performed poorly. One obvious explanation for this anomaly is that it reflects compensation for systematic risk [Johnson

(2002), Chordia and Shivakumar (2002)]. This means that winners are riskier than losers therefore command a higher expected return. The behavioural explanations for momentum suggest that it arises because investors' behaviour deviates from Bayesian [Barberis et al (1998), Daniel et al (1998)]. Both rational and behavioural explanations have received support in the literature [Zhang (2006), Chan, Jegadeesh and Lakonishok (1996), Chordia and Shivakumar (2002)]. A natural way to disentangle these explanations is to condition the performance of the momentum strategy on the initial sentiment state of investors, similarly to Cornelli et al (2006) and Livnat and Petrovic (2009). Chapter 4 discusses in more detail why behavioural theories of momentum allow an implicit link between investor sentiment and price momentum, and conditions the strategy suggested by Jegadeesh and Titman on investor sentiment.

2.3.4 Non- Expected Utility Preferences

The neoclassical model assumes that investors' choose between actions based on their expected utility by exactly figuring out the possible alternatives that are associated with each action, and weighing the utility of each alternative with its corresponding probability.

A large body of evidence has accumulated that shows decision makers systematically violate the axioms of Expected Utility. Consequently, various alternative theories of decision making have emerged that modify the traditional axioms and assumptions in an attempt to explain these observed inconsistencies. The most prominent models of non-EU preferences that have attracted the attention of financial economists are Prospect theory and models that incorporate ambiguity aversion. However, we want to stress that the literature in non-expected utility preferences is vast and many important theories are not mentioned in this section as the criterion of inclusion is the extent that the theory has been applied in Finance.

For more details on non-Expected utility models see reviews by Shoemaker (1982), Fishburn (1988) and Starmer (2000).

2.3.4.1 Prospect theory

Prospect Theory (PT), developed by Kahneman and Tversky (1979), departs from EUT in three fundamental ways. Firstly, in PT actions are not evaluated as final wealth positions, but relative to a reference point. Secondly, the shape of the utility function differs depending on whether the decision lies in the domain of gains or losses. Specifically, decision makers are risk seeking in the domain of losses and risk averse in the domain of gains, a property labelled as “loss-aversion”. And finally, in PT alternatives are not assessed based on their objective probabilities, but a transformation of them.

Prospect theory has received substantial attention from financial economists. The results suggest that models that incorporate investors with PT preferences can explain phenomena that seem anomalous to EUT-based models. For example, Barberis, Huang and Santos (2001) demonstrate that a model with loss-averse investors can explain the high mean, excess volatility and predictability of stock returns. Barberis and Huang (2008) focus on the probability weighting component of PT, and construct a model that explains the long run underperformance of Initial Public Offerings [Ritter (1991)], the private equity premium puzzle [Moskowitz and Vissing-Jorgensen (2002)], and the undpricing of distressed stocks [Campbell, Hilscher and Szilagy (2008)]. Further, Frazzini (2008) in an innovative empirical study shows that PT preferences can explain the post-earnings announcement drift [Bernard and Thomas (1992)].

2.3.4.2 Ambiguity aversion

A strong assumption of Expected Utility and the Bayesian paradigm is that investors can determine all the possible future contingencies of each asset and their associated probabilities. In reality however, probabilities are never objectively known. Nor investors receive information that explicitly refers to these probabilities. Rather, probabilities have to be determined from available information, and in certain cases they may be uncertain.⁸

Starting from the seminal work of Daniel Ellsberg (1961), it has been demonstrated that in conditions that probabilities are uncertain decision makers violate the predictions of expected utility (or subjective expected utility). In order to better illustrate the contradiction consider the typical finding in an Ellsberg-type experiment. Decision makers are asked to state their willingness to accept a gamble on the draw of a ball from two urns. The first urn contains 50% white and 50% blue balls. This urn is referred to as the “risky” urn because the decision is made in conditions of risk where the probabilities associated with the expected payoff are known with certainty. The second urn also contains white and blue balls, but in unknown proportion. Therefore, different probability distributions are possible. The decision problem is constructed as follows: in the first stage decision makers have to choose between two bets:

Bet 1: A white ball being drawn from the “risky” urn

Bet 2: A white ball being drawn from the ambiguous urn.

In the second stage choose between:

Bet 3: A blue ball being drawn from the “risky” urn

Bet 4: A blue ball being drawn from the ambiguous urn.

⁸ Savage (1954) contradicted this view and suggested that subjective probabilities always exist, thus ruled out the notion of uncertain probabilities, and extended EUT to Subjective Expected Utility.

In both cases if the subjects bet on the ball that is actually drawn they receive a payment, otherwise they receive nothing. A robust finding in the literature is that the *same* individuals choose bet 1 in the first stage and bet 3 in the second one. This is inconsistent with any kind of belief about the distribution of balls in the ambiguous urn. Since individuals choose bet 1 in the first stage they behave as if they believe that the white balls are less than 50% in the ambiguous urn. However, their betting preferences in the second stage suggest that they believe that the ambiguous urn contains less than 50% blue balls, a contradiction of the Savage axioms. Such findings demonstrate that decision makers dislike situations with uncertain or ambiguous probabilities, or that they are ambiguity averse (for a review of the evidence see Camerer and Weber 1992 and Keren and Gervitsen 1999).

Ellsberg (1961) provides a simple explanation for this phenomenon. He suggests that decision makers tend to focus on the “worst-case” probability distributions that are possible. For example, a decision maker faced with the bets in stage 1 prefers bet 1 because he thinks “what if I choose bet 2 and the frequency of white balls in the urn is too low?”, failing to acknowledge that it is equally likely that the frequency of the white balls may also be high. Ellsberg explains that decision makers in the face of ambiguity use a max-min decision criterion, whereby they maximise utility under the worst case scenario that may occur. Several models have been proposed that capture these pessimistic preferences, including multiple priors [Gilboa and Schmeidler (1989)], smooth ambiguity [Klibanoff et al (2005)] and variational preferences [Marinacci et al (2006)]. Section 3.2 in the thesis develops in more detail the model of recursive multiple priors used by Epstein and Schneider (2008) to illustrate the economics of decision making under ambiguity.

Since the seminal work of Ellsberg ambiguity aversion has been shown to be one of the most robust findings in behavioural economics. A large literature, starting with Becker and Brownson (1964) and up to the present day (Ahn, Choi, Gale and Kariv, 2009), shows

that situations that involve *ambiguity* are treated differently from those that involve risk as suggested by the original thought experiment described by Ellsberg (1961). Interestingly studies in neuroeconomics, such as Hsu, Meghna, Adolphs, Tranel and Camerer (2005), present evidence that ambiguous situations produce a unique neurological fingerprint, suggesting that ambiguity aversion is rooted in the fundamentals of human cognition.

The question that follows is whether ambiguity matters for financial markets. In traditional theories such as the Markowitz portfolio construction technique and the CAPM investors have *complete* knowledge of the distribution of expected returns, and using this information they then take positions that maximize their expected utility (see section 2.1). However, as explained in section 2.2, the probability distribution is not explicitly known but these probabilities have to be inferred from available information set. This is arguably a very difficult task because the future is highly uncertain. Imagine having to specify a single probability distribution that describes the future performance of a company. It is logical to expect that at least in some circumstances where investors do not have enough information it will be difficult to pin-down a single probability distribution. Such situations where different distributions may be possible resembles the ambiguous Urn described in this section. In such conditions assets are expected to be priced pessimistically, as shown by the large literature in decision making, contrasting the predictions of risk-based asset pricing theories.

Several studies examine the effect of ambiguity in asset markets. The general result from these studies is that assets that entail ambiguity are priced lower than what they would be if they only entailed risk, exhibiting an ambiguity premium. This ambiguity premium has been able to explain several market anomalies, such as the familiarity bias and the equity premium puzzle. The familiarity bias is the tendency of investors to prefer familiar stocks, in either their own country (the home bias), or the company in which they are employed. This tendency however, contradicts the basic prediction of neoclassical theories that investors seek

to diversify their risk. Uppal and Wang (2003) use an ambiguity aversion model and explain this anomaly. The intuition of their model is that investors are faced with less ambiguity when investing in familiar assets (either the company where they work or companies in their own country) as they have more information on which to draw and estimate their distribution of returns. Being ambiguity averse they thus prefer these investments. In addition, ambiguity aversion models explain the equity premium and the excess volatility puzzles. The equity premium puzzle, first identified by Mehra and Prescott (1985) is that historically returns on common stocks are too high, and can only be explained using unrealistically high coefficients of risk aversion. Models that incorporate ambiguity, such as Maenhout (2004), Chen and Epstein (2002) and Epstein and Schneider (2008), have been able to explain this phenomenon as they demonstrate that asset returns reflect compensation to *both* risk and ambiguity aversion. Therefore, when one uses a risk-only model to explain them the ambiguity premium is “anomalous”.

Although the theoretical literature on ambiguity is vast very few studies empirically examine whether ambiguous assets are priced pessimistically. One such study is by Sarin and Weber (1993), whereby in an experimental asset market they demonstrate that ambiguous assets are priced lower than risky ones having the same fundamental value. This finding provides empirical support to the individual experiments conducted by Becker and Brownson (1964), Bossaerts et al (2010) and Ahn, Choi, Gale and Kariv (2009), and asset pricing theories, such as Epstein and Schneider (2008). However, the literature lacks a “field-study” that examines whether the day-to-day behavior of investors exhibits traces of ambiguity aversion. The third chapter in the thesis attempts to fill this gap by examining how investors respond to ambiguity in analyst earnings forecasts. A more detailed development of the hypothesis examined and the relevant literature is given in that chapter.

The evidence reviewed in sections 2.3.4.1 and 2.3.4.2 suggest that the failure of neoclassical models to explain the functioning of financial markets may in fact lie in the assumptions of expected utility. Once we construct more realistic models of decision-making under risk we can explain many aspects of the “puzzling” behaviour of financial markets.

2.4 Limits to Arbitrage

The evidence in the previous section show that investors’ behaviour can systematically deviate from the Bayesian model. The systematic nature of these deviations rules out the notion that *if* the deviations from the Bayesian model are random, prices on average prices will be set correctly [Fama (1998)]. Rather, the evidence highlighted show that the deviations are *correlated*, and induce effects in the stock market that cannot be reconciled with the neoclassical model.

The main criticism against behavioural asset pricing, even in the presence of correlated mistakes, is the argument that behavioural biases will be more pronounced amongst individual investors, whereas institutional investors, often referred to as arbitrageurs, will behave according to the Bayesian model. Since institutional investors are fully rational, they will immediately spot any mispricing, and as they trade to exploit it they will return the price to its fundamental value. During this process capital will flow from “dumb” to “smart” investors, eventually driving the former out of the market. The arbitrage argument thus relies on two premises: *i*) that institutional investors do not exhibit any behavioural biases, and *ii*) arbitrageurs are a fully diversified army of traders with deep enough pockets that allow them to trade against a mispricing *indefinitely* until the price is returned to fundamental value.

The notion that behavioural biases are stronger amongst small investors has received support in the literature [Bonner, Walther and Young (2003), Hvidkjaer (2006), Malmendier

and Shanthikumar (2007)].⁹ However, the arbitrage argument fails in its description of risk and of arbitragers.

Firstly, the arbitrage argument omits the role of noise trader risk. Delong, Shleifer, Summers and Waldman (1990) (DSSW) explain that arbitragers do not know when the mispricing will be eliminated. This is because the sentiment of behavioural traders is unpredictable, and they can become even more optimistic or pessimistic before realizing their mistake, temporarily exacerbating the mispricing. Therefore, the model of DSSW expands the traditional notion of risk, to incorporate the unpredictable fluctuations in price caused by the sentiment of behavioural traders.

So, how does noise trader risk affect arbitrage? It is tempting to think that it makes it more profitable. That is, if the mispricing widens due to an increase in the sentiment of behavioural traders, arbitragers should engage in further arbitrage, since the opportunity to earn profits is now even greater (assuming that the asset price will eventually return to fundamental value). However, Shleifer and Vishny (1997) explain why the existence of noise trader risk, coupled with the liquidity constraints faced by arbitragers, bounds the process of arbitrage. Arbitragers are not an army of fully diversified traders; rather they are a relatively small group of highly skilled professionals employed in the various mutual and hedge funds. Secondly, they are trading *other* peoples' money. This separation of brain and capital gives rise to an agency problem, and it is precisely this agency problem that inhibits arbitrage. The argument is as follows: Arbitragers are assessed based on their performance. The holders of capital do not have the knowledge to understand the efforts of the arbitrageur to exploit a mispricing, and how this mispricing can temporarily widen. For example, suppose that the arbitrageur spots an overpriced stock, let's call it A, and sells it short, expecting to make money when the stock price declines. At the same time the arbitrageur buys a perfect substitute

⁹ Although evidence exist that professional investors are not completely immune from behavioural trading, e.g., Coval and Shumway (2005).

stock, B, so he is perfectly hedged. If in the intermediate term, noise traders become even more optimistic about A and push its price higher, whilst they become pessimistic about B and push its price lower, the arbitrageur will lose money. He will have to return stock A to its owner, thus will have to buy it at a higher price, incurring a loss, which will not be countered by his position in B. In the neoclassical model this situation is a profit opportunity, as both A and B are mispriced. Therefore the arbitrageur should ideally buy more of B and sell more of A. So long as the arbitrageur has a large enough time horizon and a large amount of capital noise trader risk does not matter. However, as Shleifer and Vishny (1997) explain, this does not occur. If the mispricing widens, the position will temporarily seem as erroneous. Capital holders, who do not have the skills to identify the mispricing, will infer that the arbitrageur is not as skilful as originally thought, and will deny or withdraw their capital. Arbitrageurs, being fearful of this scenario tend to take prudent positions, which make them more ineffective in establishing market efficiency.

Brunnermeier and Parker (2005) highlight another caveat of the arbitrage argument. They demonstrate that hedge funds (the arbitrageurs according to the neoclassical model) where *long* on technology stocks during the technology bubble at the turn of the century, managed to ride the upswing and sold before the market crashed. This evidence indicates that arbitrageurs timed the market and contributed to the mispricing. That is, they were able to identify that the average investor would continue to be optimistic about technology stocks and that prices would climb even further. Therefore, despite the fact that technology stocks with P/E ratios in the vicinity of 50 definitely seemed overpriced, they exploited the excessively optimistic sentiment of the average investor, and profited accordingly. This suggests that arbitrageurs are not always trading in terms of fundamental value, but rather by predicting the response of the average investor.

This section discussed in detail the two fundamental pillars of behaviour finance. Investors commit to *correlated* behavioural errors that are inconsistent with Bayesian updating, and that arbitrage is limited and cannot swiftly eliminate all mispricings. Thus, informational inefficiency can survive in the stock market.

The next chapters present the empirical work conducted in this thesis, explaining in detail the motivation, contribution and methodology of each chapter.

3. Ambiguity aversion and the pricing of analyst forecasts

3.1 Introduction

On May 30th 1996, Intel announced that its earnings would be sharply higher than the previous quarter, yet lower than what analysts had predicted (by 3% it turned out). This caused the price of the Intel stock to fall by 16% [Dreman (1998)]. This example illustrates that analysts' forecasts exert a powerful effect on investors' expectations, and heavily influence the allocation of wealth in the market. However, despite the value of analyst forecasts to the investment community they are noisy predictors of future earnings [Clement and Tse (2003)]. An enduring issue in finance is to identify whether investors understand the noisy properties of earnings forecasts, so that their "true" informational content is transmitted into prices efficiently [Kothari (2001)].

Rational Bayesian investors, using available information, will ex-ante identify the true distribution of forecast accuracy. Armed with this information they will respond more forcefully to forecasts that contain less noise, as shown by Equation 2.3. The result of this Bayesian updating process is that, upon impact, prices will *correctly* adjust to the substantive component of the earnings forecast, and will not demonstrate any subsequent period of adjustment.

The Bayesian approach requires that investors fully understand the behaviour of forecast accuracy, and thus the earnings process itself. However, it is unlikely that investors are able to perform this task consistently because earnings, therefore forecast accuracy, are determined in a complicated stochastic process of market-wide, industry and company-specific forces. When these processes are opaque investors may feel unable to confidently distinguish whether a particular forecast is likely to be accurate or inaccurate.

Such conditions, where the likelihoods associated with a decision are unknown, are referred to in the literature as *ambiguity*. A large literature in decision making demonstrates that people when faced with ambiguity become *pessimistic* and react as if the worst-case likelihood is the correct one [Gilboa and Schmeidler (1989), Cohen et al., (2000)]. This finding has been incorporated in models of asset prices, which predict that investors' pessimism toward ambiguous assets will induce an ambiguity premium [Chen and Epstein (2002), Maenhout (2004), Cao, Wang and Zhang (2005), Epstein and Schneider (2008), Leippold, Trojani and Vanini (2008)]. This prediction has been confirmed in experimental asset markets in which ambiguous assets consistently achieve a lower price than unambiguous ones having the same fundamental value [e.g., Camerer and Kunreuther, (1989); Sarin and Weber (1993)].

In this chapter we investigate whether ambiguity aversion affects the pricing of analyst forecasts. Specifically, we propose the pessimism-and-correction hypothesis. The initial response to ambiguous analyst revisions will be to set prices too low; downward revisions will be treated as overly bad news, while upward revisions will be treated as not such good news. Therefore, because investors will overreact to ambiguous downward forecasts and underreact to ambiguous upward ones, prices, after an initial impact period, will exhibit an upwards adjustment toward their correct levels. On the contrary, because the response of investors towards forecasts that do not involve ambiguity will not entail pessimism, we should not observe any subsequent adjustment after such forecasts.

To test the pessimism-and-correction hypothesis we must identify which factors trigger feelings of ambiguity in investors. To this end, we turn towards important theorists that have provided guidance in terms of the circumstances that can be called ambiguous. Daniel Ellsberg (1961), whose seminal paper remains the most influential, proposed that feelings of ambiguity will be related to the amount, reliability and quality of the available

information concerning the relative likelihood of events. When there is little information, or else it is unreliable or of low quality, then decision makers will feel they face a situation that is not just risky, but ambiguous. Similar *epistemic* definitions of ambiguity have been advanced by other economists and psychologists. Frisch and Baron (1998), for example, proposed that “ambiguity is uncertainty about probability, created by missing information that is relevant and could be known” (P. 1988). Einhorn and Hogarth (1985) suggest that ambiguous situations arise when the available information is vague, and does not allow one to confidently rule out alternative possibilities, while Gärdenfors and Sahlin (1982, 1983) argue that feelings of ambiguity are produced when the relevance of the available information is low. For all these authors the underlying theme is that ambiguity is negatively related to what might be called the “richness” of the information that can be used to compute relative likelihoods.

In this chapter we proxy this information richness using company size. This choice is based on previous evidence showing that in general there is less reliable information available about small companies relative to larger ones. Smaller companies receive less media coverage and they are followed by fewer analysts [Waymire (1985); Bhushan, (1989); Hong et al., (2000)], they have lower earnings quality [Imhoff (1992); Lang and Lindholm (1993), Dechow and Dichev (2002)] and higher dispersion in analyst forecasts [Diether, Malloy and Scherbina (2002), Thomas (2002)]. In sum, there is less information available concerning the earnings of smaller companies, and what information there is is harder to interpret. This, we suggest, makes forecast accuracy less predictable, and the forecast more ambiguous.

We are not the first to suggest that information about smaller companies is inherently more ambiguous. Olsen and Troughton (2000) state that because small firms entail more marginal efficiency and weaker competitive positions, it is more difficult to quantitatively

estimate their future distribution of returns. Olsen and Troughton found that professional investors agreed with them, with an overwhelming majority of respondents to a survey judging the return distributions were more ambiguous than those of large ones. As we report below, our quantitative tests of ambiguity and company size reach the same conclusion.

The chapter is divided into two parts. In the first we examine whether the size of a company for which a revision is targeted reflects ambiguity, or the degree to which forecast accuracy can be successfully predicted from available information. To perform this test we appeal to evidence that forecast accuracy can be partly predicted from publicly available information [Clement (1999), Brown (2001), Clement and Tse (2003)]. Using the comprehensive model proposed by Clement and Tse (2003) we firstly *predict* forecast accuracy, and then compare these predictions to *actual* accuracy. We hypothesized that the relationship between predicted and actual accuracy is weaker for smaller companies than for larger ones. This would indicate that for smaller companies forecast accuracy is more ambiguous, because the formation of a single posterior distribution of earnings is more difficult. Our results strongly confirm this prediction, and so support the notion that company size is a proxy for forecast ambiguity.

We then test the pessimism-and-correction hypothesis. We obtain a sample of earnings forecasts varying in their ambiguity (i.e., for companies of different size), and measure the market response to this ambiguity. We use a two-stage event study, where the first stage shows the initial impact and the second whether adjustments bring prices up to their “correct” levels. We find strong evidence for this hypothesis. Returns after ambiguous downward revisions exhibit significant reversals, while those after ambiguous upward revisions exhibit continuations. Conversely, as predicted by our hypothesis, returns after forecasts of low ambiguity do not exhibit any adjustment. These results are robust to both univariate and multivariate analysis.

In addition to the main test of pessimism-and-correction, we also investigated whether this pattern is moderated by the investor's baseline state of pessimism or optimism. We use the E/P ratio of the market to capture this baseline, because high values of this ratio indicate pessimistic investors (prices are low relative to earnings), and vice versa. Our findings show that ambiguity aversion is much stronger when the market E/P ratio is high, which suggests that the pessimism induced by ambiguous earnings forecasts is magnified when investors are generally pessimistic. This finding is in line with results from psychology showing that the degree of ambiguity aversion depends on the initial state of pessimism or optimism [Bier and Connel (1994), Pulford (2009)].

Our study contributes to the literature in three ways: Firstly, we test the generality of neoclassical theories that predict investors' responses to analyst forecasts are neither optimistic nor pessimistic [Abarbanell et al (1995)]. Consistent with the pessimism-and-correction hypothesis, we find that investors react pessimistically towards ambiguous forecast revisions, which produces price predictability.

Secondly, our results relate to the size premium [Banz (1981)]. One explanation given for this pattern is that it reflects undiversifiable covariance risk [Fama and French (1993), (1995)]. Another is that it reflects a mispricing induced by investors' behavioural biases [Baker and Wurgler (2005), Zhang (2006)]. These explanations assert that investors' estimate a single return generating process for small companies from the available information, either correctly or incorrectly. We suggest that because it is difficult to estimate a single return generating process for small companies, investors respond pessimistically to information about them, generating the size premium.

Finally, whereas previous discussions of ambiguity aversion in asset markets have either been theoretical or based on laboratory experiments [Camerer and Kunreuther (1989) and Sarin and Weber (1993), Chen and Epstein (2002)], our study provides an operational

definition of ambiguity that can be tested empirically, and tests whether ambiguity aversion affects the day-to-day behaviour of investors. Our results show that ambiguity aversion has real economic effects. These results support the view that the traditional notion of risk needs to be updated to accommodate investor's response to their feelings of ambiguity.

The remainder of the paper is organized as follows: Section 3.2 describes our methods and the sample test we conduct to examine whether our proxies relate to ambiguity. Section 3.3 defines the variables used and describes the sample. Section 3.4 presents and discusses the results, and Section 3.5 concludes.

3.2 The theory of ambiguity aversion

In order to better illustrate the theoretical underpinning of the pessimism and correction hypothesis we use the ambiguity model of information processing proposed in a recent study by Epstein and Schneider (2008) (henceforth ES).

ES model the response of a representative investor towards ambiguous information. The information is ambiguous because the signal may or may not be an accurate predictor of future performance, and the investor cannot confidently separate between the two. In the model ambiguity aversion is incorporated via the investor's utility function, namely recursive multiple priors utility [Epstein and Schneider (2003)]. This functional form incorporates Gilboa and Schmeider's (1985) maximin decision criterion whereby that agent chooses the alternative with the highest minimum utility outcome. This decision criterion implies that agents update their expectations using ambiguous information *pessimistically*. Therefore, on impact, the ambiguous information signal sends prices too low, causing an upward adjustment during an adjustment period, reflecting an ambiguity premium.

Assume a representative agent who sets prices. This agent is interested in an objective variable, θ , e.g., earnings. He has a unique and unambiguous prior over θ , defined as

$\theta \sim N(m, \sigma_\theta^2)$. This agent then receives an ambiguous information signal which yields some new information about θ defined as $S = \theta + \varepsilon$, where $\varepsilon \sim N(0, \sigma_s^2)$, $\sigma_s^2 \in [\sigma_{s,min}^2, \sigma_{s,max}^2]$. Because the signal is ambiguous its true variance (precision) is unknown. Therefore the agent forms a family of possible variances that reflect that the signal may be reliable ($\sigma_{s,min}^2$), or unreliable ($\sigma_{s,max}^2$). The agent updates in a Bayesian fashion for each different variance in the set $\sigma_s^2 \in [\sigma_{s,min}^2, \sigma_{s,max}^2]$, which results to a family of posterior distributions for the parameter θ :

$$\theta \sim N\left(m + \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_s^2}(s - m), \frac{\sigma_s^2 \sigma_\theta^2}{\sigma_s^2 + \sigma_\theta^2}\right), \quad \sigma_s^2 \in [\sigma_{s,min}^2, \sigma_{s,max}^2] \quad (3.1)$$

Being averse to ambiguity the agent makes choices according to the maximin criterion and maximizes expected utility under the worst case belief from the above set of θ posteriors. This means that the agent maximises expected utility using as an expectation the posterior distribution that entails the *lowest* θ value. This is where the asymmetric response to good and bad information derives. If the signal is good, $s > m$, the worst case scenario that entails the lower θ is when $\sigma_s^2 = \sigma_{s,max}^2$. This implies that the investor believes that the good signal is unreliable (i.e., that it has the maximum possible variance). However, the actual realization of θ will on average be higher than the investor's expectation; therefore the investor systematically underreacts towards ambiguous good information. Conversely, after bad information ($s < m$) the worst case scenario is that the information signal is totally reliable, therefore expected utility is maximised under $\sigma_s^2 = \sigma_{s,min}^2$. In this case the actual realization of θ is not as bad, therefore on average the investor overreacts towards ambiguous bad information. In both cases the expectation of the investor is the posterior that implies the lowest θ value, and it is, therefore, pessimistic.

Epstein and Schneider (2008) use this model to explain in relation to the rapid collapse of the markets in response to the 9/11 attack, followed by their gradual reversion.

Although Epstein and Schneider investigated a market-wide increase in ambiguity, we argue that a similar, albeit much smaller scale, effect operates at the firm level as well. Ambiguous news (i.e., forecast revisions) about an individual firm will lead to the same pessimism-and-correction pattern for that firm as the 9/11 attack did for the market as a whole.

3.3 Methodology and Data

3.3.1 Definition of Variables and method that ambiguity is measured

A detailed data description is provided in the next section. However, in order for the econometric methodology to be clearer to the reader we provide some information on the data used. We use analyst forecasts for quarterly earnings from the IBES database. Each forecast by each analyst for each company has a forecast release date, which in our framework is the event date. Therefore, the data in our sample are organised in a three dimensional panel related to analysts, companies and time of issuance.

Our measure of ambiguity is company size, defined as the company's market capitalization (price x shares outstanding) three days prior to the release of a forecast revision. We validated whether company size is associated with the ambiguity of forecast revisions by testing for a systematic relationship between company size and the *predictability* of forecast accuracy. As discussed above, it is not the accuracy of forecasts themselves that determine ambiguity, but the predictability of that accuracy. For instance, if an agent knows the error distribution associated with a specific forecast, she can use Bayes' rule to update her beliefs. But if the distribution is not known and different alternatives can arise, the agent will become pessimistic, as shown in Equation (3.1). In this chapter we suggest that company size relates to the extent investors can estimate the distribution of forecast accuracy.

We validate our intuition by first predicting forecast accuracy from available information using the model proposed by Clement and Tse (2003), and then examining how predicted accuracy explains actual accuracy for companies of different size. Our hypothesis is that for smaller companies actual accuracy will be more unpredictable, pointing to higher ambiguity.

In order to perform the test we firstly standardise forecast accuracy as in Clement and Tse (2003), using the following formula:

$$Accuracy_{ijtq} = \frac{AFE_{max_{jtq}} - AFE_{ijtq}}{AFE_{max_{jtq}} - AFE_{min_{jtq}}} \quad (3.2)$$

That is, the accuracy of a forecast issued by analyst i for company j in year t and quarter q is given by the absolute value of its forecast error,¹⁰ subtracted from the largest absolute error for company j in year t and quarter q made by any analyst, divided by the range of absolute errors for forecasts issued by company j in year t and quarter q . Accuracy ranges from 0 to 1, where a higher value indicates higher accuracy.

For each forecast we then computed the following variables that have been found to predict accuracy:

- 1) *Forecast horizon*: The days that separate a forecast by analyst i for company j in year t and quarter q with the corresponding earnings announcement date. The larger the distance between the two the less precise is the information used by the analyst, therefore the lower the accuracy.
- 2) *Days elapsed*: This variable is equal to the days intervening between a forecast issued by analyst i for company j in year t and quarter q , with the previous forecast issued by

¹⁰ Forecast error = |(Forecast – actual)| / actual.

any analyst for that company, year and quarter. Clement and Tse (2003) suggest that this variable captures the rate at which information for the company is flowing in. Large values indicate smaller flow thereby decreasing forecast accuracy.

- 3) *Lag accuracy*: This is a proxy for analysts' ability. It equals the absolute forecast error that analyst i has made in his last forecast for company j in quarter $q-1$. The rationale is that analysts who have been accurate in the past continue to be accurate in the future. Brown (2001) finds that lag accuracy is one of the most important predictors of current accuracy. This variable is expected to relate positively to forecast accuracy.
- 4) *Broker size*: A proxy for the reputation of the brokerage house that the analyst is employed. It equals the number of analysts employed by the brokerage house in which analyst j is employed in year t that the forecast is issued. This is expected to correlate positively with accuracy.
- 5) *Forecast frequency*: This variable captures the effort that analysts' devote to forecasting earnings for that company. It equals the amount of forecasts that analyst i has issued for company j in year t and quarter q . It is expected to correlate positively with accuracy.
- 6) *Firm experience*: Another proxy for analysts' ability. It equals the amount of time (measured in quarters) that analyst i has been covering company j . This variable also is expected to relate positively to forecast accuracy.
- 7) *Industries*: This is a proxy for the complexity of each analyst's job. It equals the number of industries (4-digit SIC) the analyst i is following in the year the forecast is issued. It is expected to relate negatively to accuracy.

8) *Companies*. Another proxy for the complexity of the task the analyst undertakes. It equals the number of companies the analyst is following in the year the forecast is issued. It is expected to relate negatively to accuracy.

These variables (except lag accuracy which is standardised as in equation 3.2) are standardised in the following way:

$$Characteristic_{ijtq} = \frac{Raw\ characteristic_{ijtq} - Raw\ characteristic_{min_{jtq}}}{Raw\ characteristic_{max_{jtq}} - Raw\ characteristic_{min_{jtq}}} \quad (3.3)$$

For example, if the analyst who issues a forecast for IBM in the first quarter of 2000 follows 3 industries, and the IBM-covering analyst who, for the same period, covers the smallest number of industries covers 1 industry, and the one who covers the largest number covers 10, then the variable in (2) for this forecast would equal $(3 - 1) / (10 - 3) = 0.28$. In this way all variables are standardised to vary between 0 and 1, with a higher value showing that the forecast is predicted to be higher for that characteristic.¹¹

Subsequently we estimate the predicted accuracy of each forecast using the following cross sectional regression in each year t (subsuming analyst and company subscripts):¹²

$$\begin{aligned} Accuracy_t = & a_t + b_{1t} Company\ experience_t + b_{2t} Broker\ size_t \\ & + b_{3t} Companies_t + b_{4t} forecast\ frequency_t + b_{5t} Days\ elapsed_t \\ & + b_{6t} Forecast\ horizon_t + b_{7t} Industries_t + b_{8t} Lag\ accuracy_t + u_t \end{aligned}$$

¹¹ Typically studies that analyze forecast errors and their association with returns do some form of standardisation because the market evaluates forecasts in a relative sense (see Brown 2001, Bonner, Walther and Young 2003, Clement and Tse 2003).

¹² We run the regression yearly, as opposed to quarterly, because accounting quarters frequently correspond to different calendar months for different companies. By running the regression in calendar time we ensure that investors can estimate expected accuracy in year t using factor loadings from $t-1$, regardless of quarter ends.

(3.4)

Using these coefficients and constant, we derive the *Expected* accuracy of forecasts issued in year $t+1$ as follows:¹³

$$\begin{aligned}
\text{Expected Accuracy}_{t+1} = & a_t + b_{1t} \text{Company experience}_{t+1} \\
& + b_{2t} \text{Broker size}_{t+1} + b_{3t} \text{Companies}_{t+1} + b_{4t} \text{forecast frequency}_{t+1} \\
& + b_{5t} \text{Days elapsed}_{t+1} + b_{6t} \text{Forecast horizon}_{t+1} + b_{7t} \text{Industries}_{t+1} \\
& + b_{8t} \text{Lag accuracy}_{t+1}
\end{aligned}
\tag{3.5}$$

Expected Accuracy is therefore an estimate of forecast accuracy using information available to the investor. As explained above, ambiguity arises when expected accuracy (Eq. 3.5) does not successfully capture actual accuracy (Eq. 3.2). To examine whether company size predicts ambiguity in this sense we regress actual accuracy on expected accuracy using dummy variables for company size:¹⁴

$$\begin{aligned}
\text{Actual Accuracy}_i = & a + a_1 \text{Middle} + a_2 \text{Small} + \theta_1 \text{Expected Accuracy}_i \\
& + \theta_2 \text{Expected Accuracy}_i * \text{Middle} + \theta_3 \text{Expected Accuracy}_i * \text{Small} + u_i
\end{aligned}
\tag{3.6}$$

Where *Middle* is a dummy that equals 1 if the size of the company belongs in the 2nd, 3rd or 4th quintile (i.e., is of intermediate size) and *Small* equals 1 if company is in the 1st quintile. The coefficient of “predictability”, θ_1 , shows how well expected accuracy (Eq. 3.5) predicts actual accuracy (Eq. 3.2) for the largest companies, those predicted to have *low ambiguity*. We predicted that as ambiguity increases, this relationship would become weaker, so that θ_2

¹³ This method of estimating expected accuracy has been used previously, e.g., Brown (2001).

¹⁴ To identify the levels of ambiguity we split our sample into quintiles based on company size. We use the breakpoints from Kenneth French’s website in month $t-1$ to assign the forecasts in month t in size deciles.

and θ_3 would be significantly negative, and that θ_3 would be less than θ_2 . That is, the smaller the company, the more accuracy is subjected to unpredictable variation that cannot be predicted from available information, and therefore the more ambiguity.

Our test of ambiguity is meaningful insofar as the market uses available information to estimate forecast accuracy, as shown in Equation 3.4. If the market does use this information it will identify a relationship, if one exists, between company size and the predictability of forecast accuracy. Although there exist plentiful evidence to suggest that investors use information cues such as those in Equation 3.4 to estimate forecast accuracy [see Clement and Tse (2003), Bonner et al (2003), Mikhail, Walther and Willis (2003); (2004)], we examine whether these findings emerge in our sample. We do this by running the following regression:

$$\begin{aligned}
 Car(-1,1)_{ijt} = & a + b_1 Company\ experience_{ijt} + b_2 Broker\ size_{ijt} \\
 & + b_3 Companies_{ijt} + b_4 forecast\ frequency_{ijt} + b_5 Days\ elapsed_{ijt} \\
 & + b_6 Forecast\ horizon_{ijt} + b_7 Industries_{ijt} + b_8 Lag\ accuracy_{ijt} + u_{ijt}
 \end{aligned} \tag{3.7}$$

CAR_{ijt} (from trading days -1 to 1 where 0 is the date the forecast is issued) is the cumulative market adjusted abnormal return following a revision from analyst i for company j in period t , and the dependent variables are the predictors of forecast accuracy used in Equation 3.4. We run the regression separately for upward and downward forecasts. This regression will show whether the market relies on predictors of accuracy to evaluate forecast quality.

3.3.2 The event study

Once we have demonstrated the validity of the ambiguity measures, we test the pessimism-and-correction hypothesis, which is that more ambiguous forecasts will be treated

more pessimistically. We group forecasts according to the size of the targeted company, and perform a two-stage event study to examine whether the average returns in each group support the existence of ambiguity aversion.

Diagram 1 below explains the intuition of the event study. The 21 trading days period after the revision, is divided into two sub-periods. The period from trading day -2 to 2, where date 0 is the date that the forecast is issued,¹⁵ is called the *impact period* and measures how expectations and prices are updated in the light of analyst forecasts revisions. The second period, from trading days 3 to 20, is called the *adjustment period* and shows how prices adjust to the initial impact. An efficient response requires that returns in the adjustment period are insignificant as the information content of the revision was fully and correctly incorporated in the asset price during the impact period.




Diagram1. The specification of the windows used.

However if returns continue on the same direction we can argue that the analyst forecast was *underused* (underreaction) and prices exhibited a drift to the equilibrium. On the contrary if returns revert the analyst forecast was *overused* (overreaction) and prices were temporarily sent to high or low and then mean reverted to equilibrium. The pessimism-and-correction

¹⁵ For date 0 the study uses the variable *estdats* in the IBES detail files.

hypothesis is that, during the impact period, investors set prices of firms with ambiguous forecast accuracy too low, and prices increase during the adjustment period to correct this initial pessimism. This means that investors are *overreacting* towards *downward* forecast revisions with ambiguous accuracy, and *underreacting* towards *upward* forecast revisions with ambiguous accuracy.

This short-run methodology has been chosen for two reasons. Firstly, Kothari and Warner (1997) and Barber and Lyon (1997) find that the statistical reliability of long-run firm specific event studies is debatable. In addition, the frequency of earnings related information, such as forecasts by other analysts and earnings announcements is so large that it is hard to disentangle the extent to which long run returns relate to one particular forecast.

Returns are risk adjusted using the modified market model.¹⁶ The market adjusted return method is an approximation of the market model where $\alpha = 0$ and $\beta = 1$, for all firms. This model has been used widely in the literature [see for example DeBondt and Thaler (1985), Fuller, Netter, and Stegemoller (2002)]. The benefit of this approach is that it does not require a pre-event estimation period in which to estimate the parameters, which will be contaminated by prior revisions that may relate to the current one, especially for closely followed larger firms skewing the results. The chapter follows the original Brown and Warner's (1985) standard event study to calculate Cumulative Average Returns (CAR) for the impact and the adjustment periods. Returns are defined as the logarithmic change of the return for the firm's security within a one day interval, R_{it} . Therefore, $AR_{it} = R_{it} - R_{mt}$ is the abnormal return for security i at day t , where R_{it} is the return of the security on that day and R_{mt} is the return of a value weighted market index including dividends. Finally, in order to calculate the Cumulative Abnormal Returns (CAR) for the two periods we sum up daily

¹⁶. The reason we choose the market model is that, as noted by Cambell, Lo and Mackinlay (1997), such statistical models suffice when one seeks to investigate the effect of an event on asset prices. This is because in smaller time horizons (i.e., less than a month) expected returns are close to 0, therefore an economic model, such as the CAPM, is not necessary.

abnormal returns for the impact and adjustment periods. In order to test the statistical significance of the CAR 's we use t -statistics estimated using the cross-sectional variation of abnormal returns. More specifically, the test statistic of the null hypothesis that the mean CAR 's ($MCAR$) is equal to zero for a sample of n firms is as follows:

$$t_{CAR} = \frac{MCAR}{\sigma_{MCAR} / \sqrt{n}} \quad (3.8)$$

Where $MCAR$ denotes the sample average, $\sigma(MCAR)$ denotes the cross-sectional sample standard deviation of the abnormal returns and n is the sample size. To test the significance of the difference between $MCAR$'s for small and large companies in the impact and adjustment period using a two sample t -test as shown below:

$$t = \frac{MCAR_{Small} - MCAR_{Large}}{\sqrt{\frac{2s^2}{N}}} \quad (3.9)$$

where the sample variance, given by:

$$s^2 = \frac{s_{small}^2 + s_{large}^2}{2} = \frac{\sum_{n=1}^N CAR_{i,small} - MCAR_{small}^2 + \sum_{n=1}^N CAR_{i,large} - ACAR_{large}^2}{2 N - 1} \quad (3.10)$$

3.3.3 Market reaction to ambiguity under general optimism and pessimism

There is evidence that the degree of ambiguity aversion is moderated when decision makers are generally optimistic [Bier and Connell (1994), Pulford et al. (2009)]. We tested for this by examining whether the magnitude of the pessimism and correction effect is also attenuated when the market is generally optimistic. The intuition is that when the market is overvalued relative to earnings (low E/P ratio) investors will be generally optimistic about the

future growth rate of the economy, and so show less ambiguity aversion than when the market is undervalued relative to earnings (high E/P ratio).

Following Conrad et. al. (2002) we proxy market optimism using the de-trended market E/P ratio. Lower values mean more optimism, because it means investors are driving up the price of assets relative to earnings. Higher values of the ratio mean prices are being driven down. The market E/P ratio is calculated by dividing total earnings by total market value for a representative sample of companies from the US market.¹⁷ To eliminate the time trend from the E/P ratio we follow Conrad et al (2002) and standardise the time series by subtracting from each month's market E/P ratio the average market E/P ratio for the preceding 12 months. Based on this standardised time series we construct deciles of (de-trended) market E/P ratio. The bottom 30% of the observations in this distribution constitutes the "good" times, and the top 30% the "bad" times. We analyze the response of the market towards low and high ambiguity forecasts in the impact and adjustment periods separately for the three market states.

3.3.4 Multivariate analysis

Our univariate analysis is complemented by regression analyses that control for other variables known to affect post-revision returns [Clement and Tse (2003)]. These are the forecast horizon (the days intervening between a forecast revision and the earnings announcement), broker size (the number of analysts employed by the brokerage house employing the analyst who issued a forecast) and revision magnitude. In addition, we control for analyst forecast bias (signed forecast error) to ensure that the pricing patterns we document are due to ambiguity aversion and not analysts systematically under predicting

¹⁷ This variable is downloaded from Thomson Financial, and is coded as TOTMKUS. Thomson report that this variable is calculated using the 1000 largest companies in the U.S.

earnings.¹⁸ Further, we include interactions of the logarithm of company size and the valuation of the market to examine whether the response to company size differs according to these variables. The multivariate regression is of the form:

$$\begin{aligned}
 CAR_{ijt} = & \alpha_0 + \beta_1 \log(size)_{ijt} + \beta_2 \log(size)_{ijt} * market\ E/P_t \\
 & + \beta_3 Market\ E/P_t + \beta_4 Forecast\ Horizon_{ijt} + \\
 & \beta_5 Broker\ size_{it} + \beta_6 Forecast\ error_{ijt} + \beta_7 Rev.magnitude_{ijt} + u_{ijt}
 \end{aligned} \tag{3.11}$$

The dependent variable, CAR_{ijt} , is the market adjusted cumulative abnormal return for company j at time t after the forecast revision issued by analyst i , while $\log(size)_{ijt}$ is the natural logarithm of company size. We conduct four regressions, for returns during the impact period and adjustment period, separately for upward and downward forecasts.

Because we are using a panel data set with repeated observations for the same companies the models in (3.6),(3.7) and (3.11) are estimated using a fixed effects regression technique that eliminates the effect of clustered residuals on the basis of year-firm combinations.

Fixed effects regressions are aimed to alleviate the fact that inevitably regression models are imperfect. They omit important covariates, potentially biasing the relevant estimates. Fixed effects regressions alleviate these misspecification problems by essentially disregarding between-cluster heterogeneity that arises due to such omitted variables, producing estimates that depend only on within-cluster variation.

Specifically, the model estimated is:

$$y_{it} = b_0 + X_{it}b + Z_i\gamma + a_i + u_{it} \tag{3.12}$$

¹⁸ See Matsumoto (2002).

Where y_{it} is the dependent variable for observation (i.e., company) i at time t , X_{it} is the time-variant regressor, Z_i is the time-invariant regressor, α_i is the unobserved individual effect, and u_{it} is the error term.¹⁹ Clearly our model demonstrates that the observations for company i in the sample are contaminated by the factor α_i . This violates the assumption of serially uncorrelated and homoscedastic residuals. Fixed effects regression eliminates this company specific component by estimating:

$$y_{it} - y_i = (X_{it} - \bar{X}_i)b + (u_{it} - \bar{u}_i) \quad (3.13)$$

where $\bar{X}_i = \frac{1}{T} \sum_{t=1}^T X_{it}$ and $\bar{u}_i = \frac{1}{T} \sum_{t=1}^T u_{it}$

Therefore the parameter of interest b is estimated from an OLS regression whereby from each observation on both the dependent and independent variables the average cluster effect is removed.²⁰

As mentioned above the estimations of models (3.6),(3.7) and (3.11) use Fixed effects regressions for clusters, i.e., α 's, based on years and companies. Year effects may arise if the application of the model for certain years differs. For example, in 2002 the SEC has approved a scheme whereby brokerages are required to have independent research and investment banking departments. If the market responded positively to this bill, we may find a stronger response to analyst forecasts subsequent that date. This would induce a year related α_i in the market response model. By including year related fixed effects, our model is robust to such between year effects.

¹⁹ Note that the procedure is named fixed effects exactly because the effect is fixed, i.e., time invariant. If the effect of thought to be time variant a random effects model should be employed.

²⁰ Numerically the estimate of b is identical to the estimate produced from the procedure outlined above if a dummy variable is included to indicate each cluster.

Clustering on the basis of companies is more obvious. Whereas all the multivariate models used apply an average relationship between the dependent and the independent variables, the relationship between y_i and x_i will most likely differ for each company in an unobserved manner. For example, assume that company i may be more transparent than company j . If transparency is not captured endogenously in the model, the residuals of these two companies will encompass an omitted transparency effect. Therefore including a company related fixed effect the model is robust to such unobserved heterogeneity.²¹

3.3.5 Sample construction and descriptive statistics

Data on all quarterly US analyst forecast revisions and actual earnings are from the IBES detail files. The sample period spans from January 1994 to December 2008.²² Data on returns and shares outstanding are from CRSP. We apply the following filters to the sample:

First, a company must have data from both CRSP and IBES. Second, if the forecast estimation date is after the earnings announcement date the observation is deleted. Third all revisions greater than 100% and all revisions which entail an error in excess of 100% of the actual earnings are removed from the sample [Capstaff et al (1995), observe that these are likely to be errors]. Fourth, all revisions equal to 0 are deleted. Finally, and consistent with other studies [Clement and Tse (2003), Clement (1999), O'Brien (1990)] only the last forecast issued by each analyst-firm pair for each quarter is retained. After these filters are applied we end up with 620,179 earnings forecast revisions, from 8,513 analysts for 6,843 companies.

Panel A of Table 3.1 presents descriptive statistics for the variables used. The means and medians for the accuracy related variables are comparable to Clement and Tse (2003).

²¹ For an in depth discussion of panel data econometrics see Petersen (2009).

²² Similar to Clement and Tse (2003) we choose 1994 as the cut-off point to ensure that the issuance date of the revision reported in the IBES files is accurate.

The average revision is -4.00%, and the average forecast error is -3.63%. These figures are consistent with prior evidence showing that on average analysts initially issue optimistic forecasts and gradually walk them down so that actual earnings meet or beat expectations [Richardson et al (2004), Matsumoto (2002)]. In terms of company size we observe that the mean size decile rank (using end of previous month breakpoints) for the firms in the sample is 6.6, which indicates the sample is slightly tilted towards large companies. This is a common finding when IBES data are used because large companies have larger analyst coverage. However the sample does include small firms, as 25% percent are below a decile rank of 4.

Panel B shows the annual frequency of analyst forecasts according to revision direction and company size. Our sample is comprised of 269,955 upward forecasts (43.4%) and 350,224 downward ones (56.6%).

Panel D1 shows that the accuracy related variables are correlated, but not to a degree that suggests problems with multicollinearity.

3.4 Results

3.4.1 Estimating forecast ambiguity

Panel A in Table 3.2 shows the output from the regression which examines whether accuracy is predictable from available information. Our results are consistent with earlier work. Accuracy increases with company experience, broker size, forecast frequency and lag accuracy, and decreases with the number of industries followed by the analyst, days elapsed and forecast horizon.

We now examine whether the market is responding to factors that predict forecast accuracy. We estimate Equation 6, whereby we regress these accuracy factors on the market

adjusted cumulative return from days -1 to 1, where 0 is the date the forecast was released. The coefficients of these factors provide information on whether the market attempts to estimate forecast accuracy when responding to the forecast.

Panel B in Table 2 shows the output from this regression.²³ Consistent with previous literature [e.g., Clement and Tse (2003), Bonner et al (2003), Mikhail, Walther and Willis (2003; 2004)], we find the market is sensitive to these accuracy factors. In terms of upward forecasts, the market responds more strongly to analysts with more experience (coefficient of 0.0008 p -value=0.024), to forecasts issued from larger brokerage houses (0.0023 with p -value <0.001), and to analysts that issue more forecasts during the quarter (forecast frequency: 0.001 with p -value 0.0148). In addition, it responds less strongly to analysts who follow many industries (-0.0017 with p -value <0.0001), and to forecasts with more days elapsed (-0.0042 with p -value<0.0001). Similar relationships are found for downward forecasts.²⁴

At this point a brief detour in the literature of forecast accuracy would be helpful to position this finding in the literature. As mentioned, because analyst earnings forecasts are very important to the investment community, an enduring issue in finance and accounting is to identify whether variations in forecast accuracy can be predicted a priori so that investment strategies that utilize analyst forecasts can be made more profitable. The early literature has produced inconclusive results, as no systematic differences in accuracy were identified (Richards 1976; Brown and Rozeff 1980; O'Brien 1987; Butler and Lang 1991,

²³ Notice that the constant term in the regression is not reported. This applies for all the regressions in the chapter. This is because the *SAS absorb* procedure (in Proc GLM) was used to eliminate year-firm effects. This procedure yields identical results as the procedure outlined in Equation 3.1, but just requires significantly less computing resources. This is important because in our case we have a very large number of year-firm observations, therefore defining each pair as a fixed effect (i.e., a dummy) substantially increase computing time, whilst producing the same output. However, for robustness, we have replicated the analysis for a subsection of the data to compare the two procedures, and the results are identical.

²⁴ However, as reported by previous studies such as Clement and Tse (2003), we find that in some cases the market responds more strongly to variables that are less predicative of accuracy. For example, the market reacts more strongly to forecasts issued earlier in the quarter (0.005 p -value <0.0001) and to analysts who follow more companies. See Clement and Tse (2003) for possible explanations for the former finding. The latter finding may indicate that analysts who follow more companies are more well known, and therefore elicit more stronger responses, all else equal.

among others). However, as research progressed several patterns began to emerge that highlight that forecast accuracy can be predicted from various factors that are observable *ex-ante*. For example, Stickel (1992) finds that analysts who are ranked as *All-Stars* are more accurate than non *All-Stars*. Clement (1999) finds that accuracy decreases with the number of companies and industries an analyst follows, and increases with the resources available to the analyst. Jacob, Lys and Neale (1999) find that forecast accuracy increases with the effort the analyst is expensing in the forecasting task. In a more recent study Clement and Tse (2003) comprehensively encompass many of these factors in a single model, and offer conclusive evidence that forecast accuracy does indeed vary with *ex-ante* forecast characteristics. The results shown in Panel A of Table 3.2 corroborate the recent findings. That is, forecast accuracy can be partly predicted from publicly available information. In addition, the results found in Panel B of table 3.2 suggest that the market is taking these factors that predict forecast accuracy into consideration when it reacts to analyst forecasts, as shown by Clement and Tse (2003), Bonner et al (2003), Mikhail, Walther and Willis (2003; 2004). These results send a clear message: investors are aware of the factors that predict forecast accuracy, and take them under consideration when they rebalance their portfolios using analyst forecast revisions.

We now estimate whether company size is associated with the predictability of forecast accuracy. Panel C in Table 3.2 shows estimates of the models in (Eq. 4) and (Eq. 5). The coefficient of predictability, θ_1 , for large companies (low ambiguity forecasts) is 0.861. Since a coefficient of 1 indicates a perfect fit, this magnitude is substantial, and suggests that the available information for larger companies can be used to predict forecast accuracy fairly accurately. The coefficient is lower for mid-sized companies (Middle = -0.0917, p -value <0.0001) and much smaller for small ones (Small = -0.1897, p -value <.0001). In other words, the predictability of forecast accuracy from expected accuracy decreases by 22% when we

move from the largest to the smallest companies. This suggests that for smaller companies the factors that the market is using to predict forecast accuracy are in significantly less relevant, therefore investors will have more difficulty distinguishing forecasts that are likely to be accurate or inaccurate. This translates to uncertainty about forecast precision for smaller companies, or, in other words, more ambiguity.

This result deserves a caveat because although we have documented the relationship between company size and forecast accuracy, size is only a *noisy* proxy of ambiguity. Therefore, we cannot be completely certain that the driving force behind our results is indeed ambiguity or some other characteristic that relates to company size. However, this is a general problem of empirical studies in finance, as we often have to measure latent quantities with observable and imperfect proxies.

3.4.2 Impact and adjustment period returns by ambiguity proxy

Table 3.3 presents market adjusted cumulative abnormal returns for equally weighted portfolios formed according to company size for the impact (-2, 2) and the adjustment (3, 20) periods. Panels A and B presents returns after, respectively, upward and downward forecasts.

We control for the magnitude of the revision, in order to measure the “quantity” of the news that each forecast releases. We measure revision magnitude as the difference between the two most recent forecasts issued by a particular analyst for a given company and quarter, scaled by the penultimate forecast. We firstly independently sort the sample into quintiles according to revision magnitude and company size.²⁵ This means that each forecast is independently put into a quintile according to firm size and revision magnitude. In order to calculate average event time returns in the impact and adjustment periods we subdivide all

²⁵ The breakpoints used to classify companies in size deciles are from Kenneth French’s website. In order to derive breakpoints for revision magnitude we sort in each month all forecasts into quintiles according to magnitude, and then use these breakpoints to assign the forecasts into revision quintiles in the next month.

forecasts in each revision magnitude quintile according to size using the original size classifications, and average the returns for each size group. This is done in order to capture the effect of company size on the pricing of analyst forecast, controlling for the magnitude of the revision. However, we also present results whereby we only sort forecasts according to company size.

Before we continue to present evidence for the ambiguity aversion hypothesis we examine whether the general finding of an unconditional underreaction toward analyst forecasts emerges in our sample [Givoly and Lakonishok (1979), Stickel (1992), Gleason and Lee (2003)]. According to this finding prices respond sluggishly to analyst forecasts, therefore continue to rise (fall) after an upward (downward) revision. It is important to verify this robust result, as we would like our sample to be representative. To check whether this holds in our sample we calculate cumulative market adjusted returns 40 days after the revision (not reported result). Consistent with previous research, we find that post revision returns drift in the direction of the revision (Givoly and Lakonishok 1979, Stickel 1991, Gleason and Lee 2003) as the average market adjusted cumulative abnormal return in the 40 day period after upward revisions is 1.89% and -2.14% after downward ones. However, our tests are different from these results because we are interested to decompose this post revision period in an impact and an adjustment period, and examine for patterns of adjustment that resemble the existence of ambiguity aversion.

We now continue to examine the pricing of analyst forecasts in the impact period. Impact periods returns also show that investors generally react more strongly to revisions issued for smaller companies. The finding is consistent with Stickel (1992) and highlights how the information set of smaller companies is poorer, which makes the market more reliant on analyst forecasts. However, we also observe that for downward forecasts the pattern is

significantly less obvious as in many cases we observe that the market is responding more strongly to larger companies (i.e., for revision magnitude quintiles 1,2 and 3 and 4).

To test the ambiguity aversion hypothesis we turn to adjustment period returns. For upward forecasts (Panel A), we observe that return continuations occur exclusively amongst smaller companies. For larger companies, regardless of the magnitude of the revision, adjustment period returns are close to zero, suggesting that the initial response has transmitted the information in analysts' forecasts into prices efficiently. However, as we move towards smaller companies returns increase monotonically in the adjustment period. For the smallest companies (Quintile 1), there is an upward adjustment in prices for all revision magnitude categories, equal to roughly 50% of impact period returns. This generates an economically and statistically significant differential in the adjustment period returns of small and large companies, which ranges from 0.50% for small revisions to 1.54% for large revisions. In short, and consistent with the pessimism and correction hypothesis, for smaller companies there is a significant upward price movement during the adjustment period.

A similar picture emerges when we consider downward revisions (Panel B). For large companies adjustment period returns for all magnitude categories are slightly negative, indicate that prices continue to drop in the adjustment period. However, as we move towards smaller companies the returns become positive, creating a significant differential between the returns of small and large companies, which ranges from 0.35% for smaller revisions to 1.35% for large revisions. This again shows that prices in the adjustment period for small companies increased substantially to push the valuation towards its correct level. This pattern is consistent with an initial overreaction which is then corrected.

Overall, the behaviour of returns in the adjustment period is consistent with the pessimism-and-correction hypothesis. In the light of ambiguous forecasts, investors overweight the pessimistic scenarios that may arise, and consequently set prices *too low*. As

the mispricing becomes apparent, prices drift upward to their “correct” levels in the adjustment period.²⁶

3.4.3 Impact and Adjustment Period Returns and the Level of the Market

In this section we examine whether the pessimism-and-correction effect is attenuated when the market is generally optimistic. We repeat the analysis in Table 3.3, by further partitioning on the E/P ratio.

Table 3.4 shows that the response of investors towards analyst forecasts does depend on whether times are good or bad. First, the market responds more strongly to upward forecasts when times are bad (i.e., the market E/P ratio is high). This is reasonable because during bad times the general price level is low and there is therefore there is more capacity for upward movements in prices. For downward forecasts impact period returns suggest that no clear relationship exists between investors’ responses and market states.

Adjustment period returns for both upward and downward forecasts show that returns are higher after ambiguous forecasts. For example, continuations after upward forecasts for smaller companies increase from 0.37% in good states, to 1.78% in bad states. The corresponding figures for downward forecasts are -0.77% to 1.86%.

Although ambiguity aversion is by far the most common attitude towards ambiguity, it has been found that in certain circumstances this pattern of behaviour is reversed. This implies ambiguity-seeking behaviour, whereby people when faced with ambiguity make optimistic choices as opposed to pessimistic ones. Various explanations have been proposed for this finding, including the competence hypothesis [Heath and Tversky (1991)], and the status-quo bias [Roca, Hogarth and Maule (2006)]. In a recent study Bracha and Brown

²⁶ One alternative possibility is that these patterns are caused by time varying expected returns. It is unlikely, however, that company risk changes so dramatically from the impact to the adjustment period, so that the continuations and reversals we document reflect compensation for bearing it.

(2009) propose that the optimistic bias can also explain ambiguity seeking behaviour.²⁷ Our findings in this section support the argument made by Bracha and Brown, as we find that in periods of general optimism, the traces of ambiguity aversion become weaker.

3.4.4 Multivariate analysis

We now examine the robustness of our results in a multivariate setting. Panels A and B of Table 3.5 show impact and adjustment stage returns for upward and downward forecasts respectively. Most results are consistent with the univariate analysis. Adjustment period returns decrease after for downward forecasts when the revision is targeted towards larger companies (-0.1581 with p -value <0.001), especially if the revision was issued during undervalued markets (coefficient of interaction term between company size and market E/P is -0.0006 with p -value <0.001).

Similar results are found for upward forecasts, as adjustment period returns decrease when the revision is targeted towards larger companies (-0.1441 with p -value <0.001). In contrast with the univariate analysis, however, we find that the interaction term between company size and market E/P is insignificant. This is perhaps a reflection of the fact that the relationship between company size and market E/P for upward forecasts is not monotonic across size groups. For example, for small companies adjustment-period returns decrease with the E/P ratio, equaling 1.78% when the E/P ratio is low, 0.89% when it is medium and 0.37% when it is high. However, for large companies no distinguishable pattern appears as the corresponding figures are -0.32%, 0.22% and 0.17%.

Impact period returns produce a less consistent picture. For upward revisions impact period returns decrease with company size, showing that the response of the market is stronger when the upward revision is issued for a small company. This is in line with the

²⁷ For studies in the optimistic bias and applications see Bracha and Brown (2009) and references therein.

univariate analysis shown in Panel A of Table 3.3, where for all revision magnitude quintiles returns are higher for smaller as opposed to larger companies. For downward revisions the coefficient of size is negative, suggesting that larger companies elicit a stronger response. This contradicts the last column in Panel B of Table 3.3, which shows that the response is *unconditionally* larger for smaller companies. However, upon closer inspection, this contradiction is captured in Table 3.3 since the unconditional effect shown in the last column only holds for the largest revisions (quintile 5 which constitutes the majority of revisions for smaller companies as shown by table 5.1), whereas for the remaining quintiles the results show that, either larger companies elicit a larger response (quintiles 2, 3 and 4) or the difference is insignificant (quintile 1). Therefore overall, the effect that dominates is that larger firms elicit a stronger response for downward forecasts.²⁸

3.4.5 Revision magnitude, company size and earnings announcements

The adjustment period in our study serves as a period in which the market adjusts its expectations in relation to its initial reaction to analyst forecasts. In this section we examine whether expectations fully adjusted during this period, or if there remains some residual pessimism that only gets corrected when the earnings announcement is made. If the adjustment period was not sufficient to fully correct earnings expectations, returns around earnings announcements will be positive and increasing with our measure of forecast ambiguity.

Table 3.6 presents the results when we link impact period returns with returns during the earnings announcement. We observe that for both downward and upward forecasts,

²⁸ Note that impact period returns do not directly relate to the hypothesis we test in this chapter. Our hypothesis is that prices will exhibit an upward movement during the adjustment period for both upward and downward forecast revisions for small companies, indicating that for small companies the initial impact was either insufficient or excessive (regardless of how it compares with the initial impact to large companies).

returns around earnings announcement increase with company size. For upward forecasts they increase from 0.37% for large companies to 1.01% for small ones, and for downward forecasts they increase from -0.19% for large companies to 0.44% for small ones. This suggests that earnings expectations did not fully adjust during the adjustment stage, therefore earnings surprises are to an extent affected by the initial pessimistic response of the market towards analyst forecasts.

3.5 Further discussion

In order to illustrate the effects of ambiguity aversion in the market, this section provides some examples where the response of the market toward ambiguous information resembled the patterns of pessimism and correction we document towards analyst forecasts.

There is evidence to suggest that aggregate market movements may also be related to ambiguity aversion. Epstein and Schneider (2008) used the theory of ambiguity aversion to explain the rapid collapse of stock prices in response to the 9/11 terrorist attacks, where the Dow Jones Industrial Average (DJIA) recorded its largest ever one-day fall on the 17th of September equal to 7.1%,²⁹ followed by the Dow's second worst week ever, with losses of 14.2%. Epstein and Schneider argued that because the 9/11 attacks were unprecedented, investors did not know how to assess their value as a "signal." They might have been a reliable predictor of a sharp drop in economic growth, in which case the fall in prices would have been justified. Or, they might have been a highly unreliable predictor, in which case they should have been disregarded and the market unaffected.

The pessimism-and-correction hypothesis predicts that investors based their initial reaction on the worst-case evaluation (that the attack was a highly reliable predictor of

²⁹ The Dow was closed for four days after the attacks.

depressed economic growth) and set prices too low. Indeed, by the end of the first week of October the DJIA returned to its pre-attack level, indicating that investors responded pessimistically, inducing an upward correction in prices during the following weeks.

A similar example may have occurred during the recent financial crisis, where the DJIA lost 56% of its value from September 2008 until March 2009. This precipitous drop might have been a rational response to news about the insolvency of major financial institutions, or it could have been an overreaction. According to the theory tested in this chapter it might have been an overreaction. Investors had no personal experience of such a crisis. Media coverage was highly equivocal, with some analysts being (relatively) optimistic and others turning to despair. Henry M. Paulson jr., in a New York Times Op-Ed from November 2008, took both positions. He started his piece with the observation that “We are going through a financial crisis more severe and unpredictable than any in our lifetimes,” and concluded it with “I am confident of success, because our economy is flexible and resilient, rooted in the entrepreneurial spirit and productivity of the American people.”³⁰ If our view is correct, such ambiguous circumstances are likely to be responded to pessimistically, and that does appear to have been what happened: At the time of this writing the DJIA has recouped almost half of that loss and continues to go up. Note that we not claiming that the drop in prices was not warranted by fundamentals (which certainly was) but rather that it may have been excessive, so markets will recoup some of that loss.

These market movements, along with the patterns of pessimism-and-correction in the pricing of analyst forecasts we have documented in this study, confirm the prediction of theoretical ambiguity models that investors price ambiguous assets pessimistically, and support the notion that the current definition of risk should be expanded to account for investors’ feelings toward ambiguity.

³⁰ The relevant report can be found online at http://www.nytimes.com/2008/11/18/opinion/18paulson.html?_r=1.

3.6 Conclusion

In this chapter we draw on a large literature in behavioural economics, which documents that people are ambiguity averse. When a single probability distribution over possible outcomes cannot be formed, people become *pessimistic* and react *as if* the worst case posterior distribution that *may* arise is the most likely to occur [Gilboa and Schmeidler (1989), Cohen et al.(2000)]. Based on this finding we formulate the pessimism-and-correction hypothesis. If the accuracy of a forecast is ambiguous and the forecast brings good news (i.e., an upward revision) the pessimistic scenario is that it is not so good news, therefore investors' underreact to it. Conversely, if it is ambiguous and brings bad news (i.e., a downward revision), the pessimistic scenario is that it is very bad news, therefore investors overreact.³¹

The decision making literature, suggests that ambiguity arises when the likelihood of future events cannot be estimated successfully from available information. Based on this definition we define as being the degree to which the market can foresee the factors that affect forecast accuracy. If forecast accuracy behaves in an unpredictable manner, investors' will not be able to confidently distinguish accurate from inaccurate forecasts, and will thus treat them as ambiguous.

We measure the predictability of forecast accuracy using company size. As noted in the literature, smaller companies have a poorer information environment compared to larger companies, and therefore can entail unpredictable variation in their earnings, making estimates of forecast accuracy less reliable.

³¹ Another case of a psychological bias affecting the response of investors towards analyst forecasts is described by Sorescu and Subrahmanyam (2006). They find that while moderate recommendations made by reliable analysts (e.g., experienced analysts in prestigious brokerages) lead to underreaction, extreme recommendations made by unreliable analysts (e.g., inexperienced analysts in less prestigious brokerages) actually produce overreaction. Sorescu and Subrahmanymam explain this in terms of a psychological theory due to Griffin and Tversky (1992), who argued that whilst decision makers overreact to the extremity of information, they underreact to its reliability.

The chapter is divided into two parts. In the first we validate that company size does indeed capture the extent that forecast accuracy is predictable. We use the comprehensive model proposed by Clement and Tse (2003) to predict forecast accuracy for companies of different size. Our results show that smaller companies entail a significant reduction in the predictability of forecast accuracy, thus support the notion that company size can index forecast ambiguity.

Next, we examine the pessimism-and-correction hypothesis. Our results confirm the predictions of this hypothesis, as we find greater continuations after upward forecasts, and greater reversals after downward ones, when these forecasts are targeted to smaller companies. When the forecasts are targeted to larger companies these patterns disappear. This suggests that for smaller companies, which entail more unpredictable forecast accuracy, the initial price was set too low, and prices rebounded to correct this initial mispricing.

These results suggest that ambiguity aversion affects the process with which analyst forecasts are transmitted into asset prices. Because investors cannot always unambiguously determine the quality of forecasts ex-ante, they update their expectations pessimistically. This response, however, leads to periods of adjustments because on average earnings are better than expected.

These results, in conjunction to an extant experimental and theoretical literature, suggest that the neoclassical definition of risk where investors have complete knowledge of the associated distributions is too simplistic. Our theories will be more descriptive if the definition of risk is expanded to account for investors feelings toward ambiguity.

Table 3.1
Descriptive statistics for analyst forecasts

Panel A presents the general characteristics of the sample. Panels B and C presents forecast frequencies in the sample partitioned by revision direction and company size in each year. To assign companies in size quintiles in month $t+1$ we use the breakpoints from Kenneth French's website in month t . Panel D shows pair wise correlation coefficients for the variables used to predict forecast accuracy. These variables are calculated as follows: forecast accuracy (defined in equation 1) on company experience (the years that analyst i has followed company j), broker size (the number of analysts employed in the brokerage of analyst i who issued the revision), companies (the number of companies followed by analyst i in year t), forecast frequency (the number of forecasts issued by analyst i for company j in year t and quarter q), days elapsed (the days that separate the forecast made by analyst i for company j in year t and quarter q with any other forecast by any analyst for the same company and period), forecast horizon (the days that separate the forecast made by analyst i for company j and the corresponding earnings announcement date), industries (the number of 4-digit SIC codes followed by analyst i in year t) and lag accuracy (the absolute error of the last forecast made by analyst i for company j in the previous quarter). These variables (except lag accuracy which is standardised according to equation 1) are standardised according to equation 2.

Panel A: Descriptive statistics

<u>Variable</u>	<u>Units of measurement</u>	<u>Mean</u>	<u>Q1</u>	<u>Median</u>	<u>Q3</u>
Company experience	Years analyst follows firm	3.06	0.83	1.92	4.17
Broker size	Number of analysts employed in brokerage in year t	59.59	23	51	93
Companies	Number of companies followed by analyst in year t	16.94	12	16	20
Forecast frequency	Number of forecasts issued by analyst	3.6	2	3	4
Forecast horizon	Days between forecast and earnings announcement	86.76	35	85	100
Industries	Number of industries followed by analyst in year t	4.63	3	4	6
Revision magnitude	(Forecast-previous forecast)/previous forecast *100	-4.00%	-11.00%	-2.30%	5.00%
Company size decile	Using end of previous month breakpoints	6.5	4	7	9
Forecast error	(forecast - actual)/actual*100	-3.63%	-11.00%	-3.03%	1.56%

Panel B: Frequencies

By revision direction:	Rev > 0	Rev < 0	Total				
	269955	350224	620179				
By size quintile:		1 (Small)	2	3	4	5 (Large)	Total
	Rev < 0	47559	55600	60094	78339	108632	350224
	Rev > 0	26734	36200	45270	62916	98835	269955

Panel C: Year by year frequencies

Year	Number of forecasts	Year	Number of forecasts
1995	26870	2002	48283
1996	28538	2003	50416
1997	31930	2004	61087
1998	37453	2005	66898
1999	41835	2006	68301
2000	39211	2007	71894
2001	42932	2008	4531
		Total	620179

Panel D: Correlation coefficients of variables to predict accuracy

	accuracy	comp. experience	broker size	companies	Forecast frequency	days elapsed	forecast horizon	industries	lag accuracy
accuracy	1	0.0056	0.0047	-0.0081	0.063	-0.037	-0.218	-0.016	0.131
		<0.0001	0.000	<0.0001	<0.0001	<0.0001	<0.0001	<0.0001	<0.0001
company experience		1	0.0461	0.184	0.112	0.119	-0.046	0.117	-0.009
			<0.0001	<0.0001	<0.0001	<0.0001	<0.0001	<0.0001	<0.0001
broker size			1	0.074	0.008	0.031	-0.006	-0.020	0.006
				<0.0001	<0.0001	<0.0001	<0.0001	<0.0001	<0.0001
companies				1	0.047	0.054	-0.009	0.514	-0.012
					<0.0001	<0.0001	<0.0001	<0.0001	<0.0001
forecast frequency					1	0.071	-0.219	0.023	0.039
						<0.0001	<0.0001	<0.0001	<0.0001
days elapsed						1	-0.076	0.047	-0.027
							<0.0001	<0.0001	<0.0001
forecast horizon							1	0.006	-0.067
								<0.0001	<0.0001
industries								1	-0.016
									<0.0001
Lag accuracy									1

Table 3.2
Ambiguity test

In Panel A we perform on OLS regression of forecast accuracy (defined in equation 1) on company experience (the years that analyst i has followed company j), broker size (the number of analysts employed in the brokerage of analyst i who issued the revision), companies (the number of companies followed by analyst i in year t), forecast frequency (the number of forecasts issued by analyst i for company j in year t and quarter q), days elapsed (the days that separate the forecast made by analyst i for company j in year t and quarter q with any other forecast by any analyst for the same company and period), forecast horizon (the days that separate the forecast made by analyst i for company j and the corresponding earnings announcement date), industries (the number of 4-digit SIC codes followed by analyst i in year t) and lag accuracy (the absolute error of the last forecast made by analyst i for company j in the previous quarter). These variables (except lag accuracy which is standardised according to equation 1) are standardised according to equation 2. A fixed effects regression is used to absorb firm-year effects. The model estimated is:

$$\begin{aligned} \text{Accuracy}_t = & a + b_1 \text{Company experience}_t + b_2 \text{Broker size}_t \\ & + b_3 \text{Companies}_t + b_4 \text{forecast frequency}_t + b_5 \text{Days elapsed}_t \\ & + b_6 \text{Forecast horizon}_t + b_7 \text{Industries}_t + b_8 \text{Lag accuracy}_t + u_t \end{aligned}$$

Panel A: The forecast accuracy model

	Parameter	Estimate	Prob-t
company experience	β_1	0.0097	<.0001
broker size	β_2	0.0112	<.0001
companies	β_3	-0.0008	0.6313
forecast frequency	β_4	0.0230	<.0001
days elapsed	β_5	-0.0351	<.0001
forecast horizon	β_6	-0.1981	<.0001
industries	β_7	-0.0049	0.0022
lag accuracy	β_8	0.0818	<.0001
Pr>F			<.0001

Table 3.2: Continued

In the top half of Panel B we regress these variables on the cumulative market adjusted return from trading days -1 to 1, where date 0 is the date the forecast was issued. The top half shows the results for upward forecasts and in the bottom half for downward. A fixed effects regression is used with dummy variables to control for fixed firm and year effects. The model estimated is:

$$\begin{aligned} \text{Car}(-1, 1)_{ijt} = & a + b_1 \text{Company experience}_{ijt} + b_2 \text{Broker size}_{ijt} \\ & + b_3 \text{Companies}_{ijt} + b_4 \text{forecast frequency}_{ijt} + b_5 \text{Days elapsed}_{ijt} \\ & + b_6 \text{Forecast horizon}_{ijt} + b_7 \text{Industries}_{ijt} + b_8 \text{Lag accuracy}_{ijt} + u_{ijt} \end{aligned}$$

Panel B: CAR (-1,1) returns on accuracy variables

REV > 0 (n=350,224)

		Estimate	p-value
company experience	β_1	0.0008	0.024
broker size	β_2	0.0023	<.0001
companies	β_3	0.0014	0.0032
forecast frequency	β_4	0.0009	0.0148
days elapsed	β_5	-0.0042	<.0001
forecast horizon	β_6	0.0047	<.0001
industries	β_7	-0.0017	<.0001
lag accuracy	β_8	0.0004	0.2082
Pr>F			<.0001

REV < 0 (n=269,055)

company experience	β_1	-0.0003	0.3694
broker size	β_2	-0.0038	<.0001
companies	β_3	-0.0016	0.0006
forecast frequency	β_4	-0.0013	0.0001
days elapsed	β_5	0.0056	<.0001
forecast horizon	β_6	-0.0095	<.0001
industries	β_7	0.0008	0.0699
lag accuracy	β_8	-0.0022	<.0001
Pr>F			<.0001

Table 3.2: Continued

In Panel C we regress actual accuracy (defined in equation 1) on predicted accuracy (defined in equation 3 and 4), with dummy variables to indicate the size quintile of the company. D1 equals to 1 if the size of the company belongs in the 2, 3 or 4th quintile, and 0 otherwise. D2 equals to 1 if the size of the company belongs in quintile 1 and 0 otherwise. A fixed effects regression is used to absorb firm-year effects. The model estimated is:

$$\text{Accuracy}_i^A = a + a_1 D1 + a_2 D2 + \theta_1 \text{Accuracy}_i^E + \theta_2 \text{Accuracy}_i^E * D1 + \theta_3 \text{Accuracy}_i^E * D2 + u_i$$

Panel C: Ambiguity test.

	Parameter	Estimate	Prob-t
Dmidsize (D1)	a1	0.1105	<.0001
Dsmall (D2)	a2	0.0606	<.0001
expected Accuracy	θ_1	0.8622	<.0001
Expected accuracy*Dmidsize	θ_2	-0.0917	<.0001
Expected accuracy*Dsmall	θ_3	-0.1897	<.0001
Pr > F			<.0001

Table 3.3:**Two-way classification of market adjusted abnormal returns for impact (-2, 2) and adjustment (3, 20) periods**

Panel A presents returns after upward and Panel B after downward forecasts. Forecast magnitude is defined as New forecast by analyst i for firm j and quarter t minus the previous forecast by the same analyst for the same company, divided by the previous forecast. Size is defined as the market capitalization of company j (price 3 days prior to the forecast \times shares outstanding). To assign companies in size quintiles in month $t+1$ we use the breakpoints from Kenneth French's website in month t . The table provides p -values for the significance of the differentials that are adjusted for unequal variances.

(-2,2)	Panel A: REV > 0						Panel B: REV < 0					
	Revision Magnitude						Revision Magnitude					
Company size	1(Small)	2	3	4	5(Large)	All	1(Small)	2	3	4	5 (Large)	All
1(Large)	0.72	1.04	1.21	1.53	1.68	1.09	-0.39	-0.99	-1.78	-2.81	-3.80	-1.65
2	1.27	1.61	2.00	2.60	2.48	1.9	-0.33	-0.93	-1.72	-2.50	-3.73	-1.86
3	1.65	1.94	2.43	2.70	3.09	2.32	-0.25	-0.90	-1.85	-3.04	-4.95	-2.48
4	1.64	2.33	2.64	3.14	4.01	2.8	-0.33	-0.79	-2.01	-3.04	-5.09	-2.86
5(Small)	1.27	1.88	2.30	2.55	3.82	2.59	-0.40	-0.81	-1.39	-2.55	-4.68	-2.93
Dif 5-1	0.55	0.84	1.10	1.02	2.14	1.50	-0.01	0.18	0.39	0.27	-0.88	-1.28
p-value	0.060	<0.0001	<0.0001	<0.0001	<0.0001	<.0001	0.970	0.17	0.00	0.03	<0.0001	<.0001
(3,20)												
1(Large))	-0.03	0.11	0.12	0.14	0.03	0.06	-0.18	-0.27	-0.33	-0.16	-0.25	-0.23
2	0.16	0.10	0.17	0.34	0.52	0.22	0.06	-0.15	0.26	0.03	0.55	0.15
3	-0.03	0.11	0.34	0.47	0.80	0.31	-0.17	0.12	-0.01	0.40	0.37	0.17
4	0.32	0.21	0.59	0.49	1.04	0.53	0.34	0.48	0.44	0.55	0.69	0.54
5(Small)	0.50	0.83	0.93	1.23	1.54	1.1	0.17	0.39	0.60	0.54	1.09	0.74
Dif 5-1	0.53	0.72	0.81	1.09	1.50	1.04	0.35	0.66	0.92	0.70	1.35	0.97
p-value	0.070	<0.0001	<0.0001	<0.0001	<0.0001	<.0001	0.180	0.00	<0.0001	<0.0001	<0.0001	<.0001

Table 3.4:**Two-way classification of market adjusted abnormal returns for impact and adjustment periods by company size and market P/E ratio**

Panel A presents returns after upward forecasts and Panel B after downward. Size is defined as the market capitalization of company j (price 3 days prior to the forecast x shares outstanding). To assign companies in size quintiles in month $t+1$ we use the breakpoints from Kenneth French's website in month t . Market E/P is calculated as the ratio of total earnings divided by total market values of the 100 largest US companies. This variable is readily available from Thomson. In order to de-trend the time series each months market E/P is subtracted from the average market P/E for the preceding 12 months. The bottom 30% of this standardised time series is classified as "low" and the top 30% as "high". The table provides p-values for the significance of the differentials calculated using standard errors adjusted for unequal variances.

company size	Market E/P ratio									
	Panel A: Rev > 0					Panel B: Rev < 0				
	1(High)	2	3(Low)	Dif 3-1	P value	1(High)	2	3(Low)	Dif 3-1	P value
(-2,2)										
1(Large)	1.29	1.12	0.88	-0.41	<.0001	-1.61	-1.60	-1.78	-0.17	0.01
2	2.36	1.80	1.56	-0.80	<.0001	-2.09	-1.64	-1.98	0.11	0.18
3	2.72	2.38	1.77	-0.95	<.0001	-2.39	-2.24	-3.02	-0.63	<.0001
4	2.84	3.20	1.97	-0.87	<.0001	-3.09	-2.56	-3.14	-0.05	0.75
5(Small)	2.92	2.57	1.96	-0.96	<.0001	-3.46	-2.51	-3.10	0.36	0.03
Dif 5-1	1.63	1.45	1.08			-1.85	-0.91	-1.32		
p-value	<.0001	<.0001	<.0001			<.0001	<.0001	<.0001		
(3,20)										
1(Large)	-0.32	0.22	0.17	-0.50	<.0001	0.14	-0.17	-0.70	-0.84	<.0001
2	0.12	0.32	0.18	0.06	0.56	1.01	0.00	-0.65	-1.66	<.0001
3	0.66	0.33	-0.11	-0.77	<.0001	0.94	0.00	-0.53	-1.47	<.0001
4	0.54	0.80	0.00	-0.54	0.002	1.16	0.71	-0.56	-1.72	<.0001
5(Small)	1.78	0.89	0.37	-1.41	<.0001	1.86	0.62	-0.77	-2.63	<.0001
Dif 5-1	2.10	0.67	0.20			1.72	0.79	-0.07		
p-value	<.0001	<.0001	0.40			<.0001	<.0001	0.66		

Table 3.5
Multivariable regressions

Impact (-2,2) and adjustment (3,20) period returns for upward (Panel A) and downward (Panel B) forecasts are regressed on the inverse of the logarithm of company size (price 3 days prior to the forecast x shares outstanding) with interactions with the de-trended market P/E ratio (the ratio of total market values divided by total earnings of the 100 largest US companies, which is available from Thomson) and revision magnitude ((the difference between the two most recent forecasts issued by analyst *i* for company *j* during period *t*, scaled by the penultimate forecast), the market E/P ratio, forecast horizon (the days that separate a particular forecast with the corresponding earnings announcement date), broker size (the number of analysts employed in the brokerage during the year the analyst *j* has issued a forecast), forecast error (the difference between forecasted value and actual value, scaled by actual value) and revision magnitude. A fixed effects regression is used to absorb firm-year effects. The model estimated is:

$$CAR_{ijt} = \alpha_0 + \beta_1 \log(\text{size})_{ijt} + \beta_2 \log(\text{size})_{ijt} * \text{market E/P}_t + \beta_3 \text{Market P/E}_t + \beta_4 \text{Forecast Horizon}_{ijt} + \beta_5 \text{Broker size}_t + \beta_6 \text{Forecast error}_{ijt} + \beta_7 \text{Rev.magnitude}_{ijt} + u_{ijt}$$

Variable	Parameter	Impact period		Adjustment period	
		Estimate	Pr > t	Estimate	Pr > t
Panel A: Rev > 0					
log(size)	b1	-0.0898	<.0001	-0.1441	<.0001
log(size)*market P/E	b2	-0.0006	<.0001	0.0000	0.5954
Market P/E	b3	0.0062	<.0001	0.0018	0.0042
Forecast Horizon	b4	0.0000	0.9827	0.0000	0.0017
Broker size	b5	0.0000	<.0001	0.0000	0.1265
Forecast error	b6	0.0004	0.6555	-0.0208	<.0001
Revision magnitude	b7	0.0279	<.0001	0.0042	0.0041
Pr > F			<.0001		<.0001
Panel B: Rev < 0					
log(size)	b1	-0.1076	<.0001	-0.1581	<.0001
log(size)*market P/E	b2	-0.0012	<.0001	-0.0006	<.0001
Market P/E	b3	0.0138	<.0001	0.0067	<.0001
Forecast Horizon	b4	0.0000	<.0001	0.0000	<.0001
Broker size	b5	0.0000	<.0001	0.0000	0.7025
Forecast error	b6	-0.0096	<.0001	-0.0199	<.0001
Revision magnitude	b7	-0.0799	<.0001	-0.0113	<.0001
Pr > F			<.0001		<.0001

Table 3.6
Returns around earnings announcements

Returns are the cumulative market adjusted returns for the period -5,0, where 0 is the announcement date. Size is defined as the market capitalization of company j (price 3 days prior to the forecast \times shares outstanding). To assign companies in size quintiles in month $t+1$ we use the breakpoints from Kenneth French's website in month t . The table provides p-values for the return differential between the extreme portfolios which is adjusted for unequal variances.

Panel A: Rev > 0		
1/company size	Impact	earnings announcement
1=Low	1.09	0.37
2	1.90	0.53
3	2.32	0.68
4	2.80	0.79
5=High	2.59	1.01
 Dif (5-1)	 1.5	 0.64
prob-t	<.0001	<.0001
 Panel B: Rev < 0		
1=Low	-1.65	-0.19
2	-1.86	0.15
3	-2.48	0.39
4	-2.86	0.34
5=High	-2.93	0.44
 Dif (5-1)	 -1.28	 0.63
prob-t	<.0001	<.0001

4. Investor Sentiment and Price Momentum

4.1 Introduction

An extensive body of literature documents that short-run stock returns are positively correlated with past returns, a phenomenon referred to as price momentum [Jegadeesh and Titman (1993, 2001), Chan, Jegadeesh and Lakonishok (1996)]. This return pattern is found to be robust in different markets [Rouwenhorst (1999), Doukas and McKnight (2002)], and different asset classes [Asness, Moskowitz and Pedersen (2008)]. The explanations for momentum proposed in the literature fall into three general categories: theories of market frictions [Hong and Stein (1999)], theories of time-varying expected returns [Johnson (2002)], and behavioural theories of market inefficiency [Daniel, Hirshleifer and Subrahmanyam (1998)].

Theories of market frictions suggest that the positive autocorrelation in returns arises because the market responds to information with a lag. This is because information is not instantaneously available to all market participants, therefore, it is sluggishly transmitted into prices. This induces positive correlation in returns, which appears anomalous in the neoclassical model which assumes that prices immediately adjust to new information. The explanation of time-varying expected returns suggests that the spread between winner and loser portfolios arises due to some form of undiversifiable covariance risk. That is, for some reason past winners are riskier relative to past losers, inducing a positive and significant spread. For example Johnson (2002) suggests that growth rate risk rises with actual growth in prices, therefore past winners are riskier. Lastly, theories of market inefficiency suggest that momentum reflects a mispricing. That is, initially information was not transferred into prices efficiently, which induces an adjustment in prices making prices predictable. These theories of market inefficiency can be divided into behavioural theories of under and over reaction.

Barberis, Shleifer and Vishny (1998) suggest that investors' underreact to information, therefore underpricing the winners and overpricing the losers, which induces momentum in the next period. Daniel et al (1998) suggest that momentum is actually overreaction induced by overconfidence and biased self attribution. The difference between these two classes of theories is that the overreaction story predicts that momentum profits will revert in the long run.

These explanations for momentum have received support from empirical studies, creating tension in the literature. For example, Hong, Lim and Stein (2000) show that, controlling for firm size, momentum profits are decreasing in analyst coverage, thus support the notion that momentum is caused by slow information diffusion. Chordia and Shivakumar (2002) find that momentum profits are entirely predictable from a set of macroeconomic variables, proposing a rational explanation of momentum. Cooper, Gutierrez and Martin (2004) find that momentum returns are entirely captured by lagged market returns, and suggest a behavioural explanation of momentum.³² Resolving this tension is important. This chapter examines whether behavioural theories provide an adequate explanation for momentum profits by examining the relationship between momentum profits and investor sentiment.

Sentiment, broadly defined, refers to whether an individual, for whatever *extraneous* reason, feels excessively optimistic or pessimistic about a situation. A large body of the psychology literature finds that peoples' *current* sentiment affects their judgment of future events. For example, Johnson and Tversky (1983) show that people that read sad newspaper articles subsequently view various causes of death, such as disease etc., as more likely than people who read pleasant newspaper articles. In general, the evidence from experimental

³² For further discussions on the origins of momentum see Conrad and Kaul (1998), Moskowitz and Grinblatt (1999), Grundy and Martin (2001) and Grinblatt and Han (2005).

psychology shows that people with positive sentiment make optimistic judgments and choices, whereas people with negative sentiment make pessimistic ones [Bower (1981, 1991); Arkes, Herren, and Isen (1988); Wright and Bower (1992); among others].

Investor sentiment is relevant to the behavioral theory proposed by Daniel et al. (1998). In their model, overconfident investors analyze information and form overly extreme expectations about the future prospects of companies. Investors tend to disregard news that is contrary to these expectations, which can create momentum, pushing prices above fundamental values, ultimately leading to reversals. The arguments of Daniel et al. (1998) indicate that when investors are generally optimistic, they will ignore new contradictory information about firms due to their self-attribution bias. This activity, in turn, will positively impact stock prices, and amplify the momentum effect. Short-selling constraints will prevent arbitrage from correcting prices. A symmetric effect need not obtain in the case where investors are pessimistic, because even if pessimists ignore good news about a firm, arbitrage forces will incent them to buy stocks and correct prices accordingly. In this study, we examine the short- and long-run performance of momentum portfolios that are conditional on investor sentiment. We predict that when investors are optimistic, short-run momentum profits will be higher, and will exhibit long-run reversals.

The studies by Chordia and Shivakumar (2002) and Cooper et al. (2004) are related to our analysis. Cooper et al. (2004) suggest that investors' behavioral biases will be more accentuated after market gains, and show that momentum is profitable only after increases. They interpret this finding as supportive of behavioral explanations for momentum. Our study corroborates this evidence by partitioning momentum profits on investor sentiment, a potentially more direct proxy of investors' propensity to form erroneous beliefs. We show that sentiment has incremental power to explain momentum-induced profits even after accounting for market returns. Chordia and Shivakumar (2002) show that momentum profits

are only significant in periods in which the economy is expanding, and put forward a rational explanation of momentum. However, these authors are careful to point out that their findings are entirely consistent with a behavioral story where investors generate momentum during market expansions because they are excessively optimistic.³³ This is precisely the avenue we pursue in our study. We condition momentum profits on investor sentiment, and predict that momentum profits will be higher when investors are optimistic, and will eventually lead to long-term reversals as this optimism is reversed.

To ensure that our CB Index is free of macroeconomic influences, our investigation uses an orthogonal version of the index, which is obtained by regressing the CB Index on a set of macroeconomic variables. The variables include growth in industrial production, real growth in durable, non-durable, and services consumption, growth in employment, and a National Bureau of Economic Research (NBER) recession indicator. Furthermore, we examine the sensitivity of our results to an alternative index for investor sentiment constructed by Baker and Wurgler (2006, 2007) (BW).

We show that when investor sentiment is optimistic, the six-month momentum strategy yields significant profits, equal to an average monthly return of 2.00%. However, when investor sentiment is pessimistic, momentum profits decrease dramatically to an *insignificant* monthly average of 0.34%. We also find that investor sentiment provides an important link between short-run continuation and long-run stock price reversal. We examine the long-run behavior of optimistic and pessimistic momentum portfolios six years after portfolio formation, and find that momentum profits revert after optimistic periods, with a substantial average monthly loss of -0.56%, whereas momentum profits after pessimistic

³³ Chordia and Shivakumar (2002) suggest that the challenge to this rationale would be to provide an explanation of why investors misinterpret market-wide information and become overly optimistic, misreacting to company-specific information. Investor sentiment provides such an explanation, since the general finding is that optimism that is unrelated to the decision at hand, i.e., optimism related to the state of the economy and not the individual company, can alter the choice made.

periods do not.

Our tests help disentangle rational from behavioural explanations of momentum. We note that rational theories do not allow a role for investor sentiment in causing momentum or reversals. Further, we show that our results are robust to different size- and volume-sorted portfolios, alternative proxies for investor sentiment, the CAPM with conditional and unconditional betas, Fama-French (1993) risk adjustments, and controls for microstructure biases. Since our findings do not have any obvious rational explanation based on frictions or risk, our study indicates that behavioural theories are a more appropriate fit for the data.

Secondly, our study is related to the sentiment literature, which has produced important evidence that suggests that sentiment is priced. This has led several authors to explore the relationship between investor sentiment and various stock market anomalies. Along these lines, investor sentiment has been linked to the post earnings announcement drift [Livnat and Petrovic (2008)], fund flows and the value effect [Frazzini and Lamont (2008)], corporate disclosure [Bergman and Roychowdhury (2008)], IPOs [Cornelli, Goldreich, and Ljungqvist (2006)], and the size effect [Baker and Wurgler (2006, 2007)]. Our study extends this literature by analyzing the relationship between investor sentiment and momentum, an important stock market anomaly.

Further, our results expand the literature that views momentum as a behavioural phenomenon. Zhang (2006) shows that because investors have a tendency to underreact, they underreact even more generating higher momentum when companies are perceived to be subject to higher information uncertainty. Cooper, Gutierrez and Hameed (2004) investigate the relationship between investor overconfidence and momentum, using UP market conditions to proxy for investor overconfidence resulting from self-attribution bias. They confirm that higher overconfidence leads to higher momentum returns [see also Daniel and

Titman 1999]. Our study extends this literature by investigating whether investor sentiment, a purely behavioural attribute, affects momentum profits.

Our results also disentangle theories of underreaction [Barberis et al (1998)] and theories of continued overreaction [Daniel et al (1998), Hong and Stein (1999)], which have received support in the literature. For example, Cooper et al (2004) and Lee et al (2000) provide evidence that in certain conditions momentum profits revert, supporting the overreaction theory, whereas Chan, Jegadeesh and Lakonishok (1996) show that momentum does not revert, supporting the underreaction theory. Our study shows that momentum profits revert after optimistic periods, which provides strong support to the overreaction theories of momentum.

This chapter is organized as follows. Section 1 describes the data and the empirical methodology. Section 2 presents the results, along with a discussion of the sensitivity analysis and robustness checks. Section 3 concludes the chapter.

4.2 Data and Methodology

We use all common stocks (share codes 10 and 11) listed in the New York and American Stock Exchanges (NYSE and AMEX respectively) from the Center for Research in Security Prices (CRSP) monthly file. The sample time period is from February 1967 to December 2008, for which the monthly CB Index is available.

We construct momentum portfolios using the methodology of Jegadeesh and Titman (1993). In each month t , we sort all stocks on their returns for the past J months. Based on these rankings, ten equally weighted portfolios are formed. The top decile is called the “losers” portfolio, and the bottom decile the “winners” portfolio. Every month, the strategy takes a long position in the winner portfolio and a short position in the loser portfolio, held for K months. We construct overlapping portfolios to increase the power of our tests.

Specifically, we close the position initiated in month $t-K$ in both the winner and loser portfolios, and take a new position using the winners and losers of month t . Therefore, in each month, we revise $1/K$ of the stocks in the winner and loser portfolios, and carry over the rest from the previous month.³⁴ In order to avoid microstructure biases, we allow one month between the end of the formation period and the beginning of the holding period, and delete all stocks that are priced less than one dollar at the beginning of the holding period.

As mentioned earlier, for the main part of our analysis we measure investor sentiment using the monthly time series of consumer confidence sentiment constructed by the CB. This survey began on a bimonthly basis in 1967 and turned into a monthly series in 1977.³⁵ The CB questionnaire is sent to 5,000 randomly selected households in the United States, and asks participants five questions about their outlook for the economy.³⁶ The scores for each question are calculated as the number of favorable replies, divided by the sum of favorable and unfavorable replies. The scores on the five questions are amalgamated to form the overall Consumer Confidence Index. The Index is one of the ten leading economic indicators published by the CB, and has been used in studies to predict household spending activity [Acemoglu and Scott (1994); Ludvigson (2004)]. Further, such measures of consumer confidence are positively related to investor optimism [Fisher and Statman (2002)], and have been used as proxies for investor sentiment [e.g., Lemmon and Portniaguina (2006)].

In order to purge the effects of macroeconomic conditions from the CB Index, we regress this monthly index on six macroeconomic indicators: growth in industrial production,

³⁴ For example, for the six-month formation-holding period strategy ($J, K=6$), in each month $t+1$, the winner portfolio is comprised of $1/6$ (winners from $t-1$) + $1/6$ (winners from $t-2$) + ... + $1/6$ (winners from $t-6$), and correspondingly for the loser portfolio. Note that month t is skipped.

³⁵ For the period that the index is available on a bimonthly basis, we follow Qiu and Welch (2006) in using linear interpolation to obtain monthly observations.

³⁶ The questions are the following: 1) How would you rate present general business conditions in your area? 2) What would you say about available jobs in your area right now? 3) Six months from now, do you think that the business conditions in your area will be better, same or worse? 4) Six months from now, do you think there will be more, same, or fewer jobs available in your area? 5) Would you guess your total family income to be higher, same, or lower 6 months from now?

real growth in durable consumption, non-durable consumption, services consumption, growth in employment, and an NBER recession indicator, and use the residuals from this regression as the sentiment proxy.³⁷

To identify whether a particular formation period is optimistic or pessimistic, we calculate a weighted rolling average of the sentiment level for the three months prior to the end of the formation period. We give a weight of three to sentiment in the prior month, two to the one in the month prior to that and one to the month three months prior to the current month.^{38,39} In order to ensure that our analysis is not sensitive to the definition of sentiment states, we report results using three different classifications of optimistic and pessimistic investor sentiment states. In the first specification a formation period is classified as optimistic (pessimistic) if the three-month rolling average ending in month t belongs in the top (bottom) 30% of the three-month rolling average sentiment time series, whereas for the second and third specifications we consider 20% and 40% cut-off points.

Because we form overlapping portfolios, in each holding period month we hold stocks from K different formation periods, across which sentiment can differ. In order to calculate the average sentiment in these K formation periods, we first calculate whether each of these K formation periods was optimistic or pessimistic as explained above, and then tally how many

³⁷ This sentiment indicator is also used by McLean and Zhao (2009).

³⁸ The CB Index for month $t-1$ is made publicly available from the beginning of month t . Thus, to make sure that all the information we use is available upon portfolio construction, we classify the momentum portfolio formed at the end of month t as optimistic or pessimistic using the weighted average of the residual sentiment from the three previous months as follows: $3/6*\text{residual}(t) + 2/6*\text{residual}(t-1) + 1/6*\text{residual}(t-2)$. This weighting scheme is chosen in order to assign more weight on the most recent sentiment observation when we predict momentum profits, and is similar to the one used in Lakonishok, Shleifer and Vishny (1994, p. 1550). However, our main results remain unchanged when we use a simple arithmetic average.

³⁹ Since sentiment is announced with a one-month delay, the use of residuals from month t , $t-1$ and $t-2$ to calculate the rolling sentiment measure actually corresponds to sentiment during months $t-1$, $t-2$, and $t-3$. We also consider alternative sentiment specifications based on two and four month lags and find that our results continue to hold. These results are reported later in the paper.

were optimistic or pessimistic.⁴⁰ If, all the K formation periods were classified as optimistic (pessimistic) the particular holding period month is classified as optimistic (pessimistic), with the rest being the “mild” sentiment months.

To test whether momentum profits in each sentiment state are equal to zero, and regress the time series of average monthly momentum profits on an *OPTIMISTIC* sentiment dummy variable, a *MILD* sentiment dummy variable and a *PESSIMISTIC* sentiment dummy variable, with no intercept. To test if mean profits in optimistic sentiment periods are different from profits in pessimistic sentiment periods, we regress average monthly momentum profits on a *MILD* sentiment dummy variable and an *OPTIMISTIC* sentiment dummy variable with a constant. This approach, which is similar to that of Cooper et al (2004), helps preserve the full-time series of returns, and allows us to estimate t -statistics that are robust to autocorrelation and heteroskedasticity using Newey and West (1987) standard errors.

We also calculate the long-run performance of the momentum portfolios, focusing on the six-month formation/holding period strategy. We follow the methodology employed by Jegadeesh and Titman (2001), whereby for each momentum portfolio constructed, we define an event time that is equal to 13 months following the initial formation date.⁴¹ After this event date, we hold the portfolio for six years, and test whether portfolios formed in optimistic formation periods behave differently from those formed after pessimistic formation periods.

Table 4.1 presents descriptive statistics for our sentiment index. Panel A is based on

⁴⁰ For example, assuming $K=6$, in June 1980 we hold stocks selected from six ranking periods ending in May, April, March, February, and January. For each of the six ranking periods, we calculate the sentiment level in the previous three months, and classify each formation period as being high, mild, or low sentiment.

⁴¹ Thua, the portfolio held in June 1980, for instance, was initiated in November 1979 (skipping December). This portfolio is based on overlapping returns, thus it is an equally-weighted portfolio of the positions initiated in January, February, March, April, and June. For this portfolio, the post-holding period starts in January 1981, after which we continue to hold the same portfolio using the equally-weighted structure for a period of six years.

the raw data of consumer confidence provided by the CB. Panel B reports the three-month rolling average using the residuals from regressing the raw CB data on a set of macroeconomic variables. The raw CB Index, as shown in Figure 4.1, rises during the late 1960s, mid 1980s, and late 1990s, and falls during the 1970s and early 1990s. These patterns are in line with the evidence for investor sentiment discussed by Baker and Wurgler (2006). The fall in sentiment for the period 2006-2008 seems to be a reflection of the early signs of the current recession. As shown in Figure 1, the 3-month rolling weighted average of this residual, which is the sentiment measure used in our main analysis, tracks the raw CB index closely (i.e., shows an upward trend when the index is rising and vice versa).⁴²

A robust finding in the literature is that investor sentiment is reflected in the size premium [Lee, Shleifer, and Thaler (1991); Baker and Wurgler (2006, 2007); Lemmon and Portniaguina (2006)]. The interpretation given to this finding is that optimistic investors are drawn to small stocks, thereby reducing the size premium in the following period. In order to validate our sentiment proxy, we test whether it captures this negative relationship with the size premium. Specifically, we regress the three-month average of residual sentiment ending in month t on the return of the Small minus Big portfolio (SMB) in month $t+1$ and a constant.⁴³ Indeed, as expected, we obtain a coefficient of -0.023 (t -value = -3.16), which corroborates our proxy as a sentiment index.

⁴² The fact that the orthogonalization does not materially affect the behaviour of the index is in line with the findings of Baker and Wurgler (2006).

⁴³ We thank Kenneth French for making the SMB data available on his website (<http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/>).

4.3 The Empirical Evidence on Momentum Profits across Sentiment States

4.3.1 Investor Sentiment and Short-Run Momentum Profits

Our first empirical test examines the profitability of the momentum strategy conditioning on pessimistic and optimistic investor sentiment states. Table 2 presents the results for strategies that are based on a six-month ranking period (J) and holding periods (K) of three, six, and twelve months sorted by investor sentiment. In Panel A (B, C) pessimistic sentiment is defined as the bottom 30% (20%, 40%) of the rolling average sentiment time series.

The unconditional momentum strategy for the period 1966-2008, based on $J, K=6$, yields an average monthly profit of 1.38% (unreported result). This figure is comparable with studies of momentum for analogous time periods [Lee and Swaminathan (2000); Jegadeesh and Titman (2001)]. Note, however, that momentum profits are extremely sensitive to investor sentiment. In Panel A of Table 4.2, the six-month strategy ($J=6, K=6$) shows that the average monthly profits in optimistic periods are highly significant, at an average of 2.00% per month. These profits decrease to 1.46% per month in mild sentiment months, and shrink to a statistically insignificant monthly average of 0.34% in pessimistic sentiment months. When the holding period is extended to twelve months ($J=6, K=12$), average monthly profits in optimistic (mild) periods are 1.27% (0.85%), and while they decline to 0.09% in pessimistic periods. Similarly, when the holding period is condensed to three months ($J=6, K=3$), average monthly profits in optimistic (mild) periods are 2.28% (1.44%), declining to an insignificant 0.28% in pessimistic periods. Comparable results are shown in Panels B and C of Table 4.2, confirming that the choice of cut-off point for optimistic and pessimistic sentiment does not materially affect our conclusions.⁴⁴

⁴⁴ In Panel B3 ($J=6, K=12$ with 20%-20% sentiment states) if we identify sentiment states using the method described in Table 2 we end up with a very small number of observations in the optimistic and pessimistic

The finding that momentum profits are significant in “mild” sentiment months suggests that *on average* investors beliefs are optimistic, which is in line with extant empirical evidence [e.g., Weinstein (1980), Slovic (2000), Puri and Robinson (2007)]. Consequently, momentum profits arise in “normal” sentiment periods, increasing (decreasing) dramatically when sentiment becomes more optimistic (pessimistic). As stated earlier, these results corroborate the analysis of Chordia and Shivakumar (2002) and Cooper, Gutierrez, and Hameed (2004), who respectively find that momentum profits vary significantly according to whether the market has been rising or falling or whether the economy has been expanding or contracting. Going further, however, our analysis explicitly links the time series of momentum profits to investor sentiment.

Another interesting result that emerges from Table 4.2 is that returns of all momentum portfolios during pessimistic periods are *higher* than those in optimistic periods across all holding period horizons. This result is consistent with previous findings [Baker and Wurgler (2006, 2007)], suggesting that investors tend to overestimate the likelihood of *negative* events when they are pessimistic, setting prices lower. Furthermore, Table 2 demonstrates that higher profits due to momentum strategies in optimistic periods arise primarily because loser stocks exhibit higher momentum than winner stocks during pessimistic periods. This is consistent with the notion that optimistic investors disregard negative information about loser stocks and short-selling constraints limit arbitrage activity. This negative information, however, is eventually incorporated into prices, causing underperformance in the loser portfolio, generating momentum profits. A symmetric effect does not obtain during pessimistic periods because investors who ignore positive information during pessimistic

sentiment group. For this reason we identify optimistic and pessimistic states as follows: we assign a score to each formation period depending on its sentiment. If the particular formation period is optimistic we assign it a score of 2, a score of 1 if it is mild and a score of 0 if it is pessimistic. When we average returns for the particular holding period month we also average the sentiment score of that holding period (from the K formation periods). If the average sentiment score is > 1.5 (≤ 0.5) the month is identified as optimistic (pessimistic). The rest are the mild sentiment months.

periods have their bias countervailed by arbitrage buyers.

Our results suggest that the momentum trading style is *not* a risk-free arbitrage opportunity, as the returns of the winner and the loser portfolios do not preserve their spread across both optimistic and pessimistic sentiment states. Significant profits obtain, however, when the momentum strategy is implemented only after optimistic periods.

4.3.2 Is the Effect of Investor Sentiment on Momentum Profits Robust?

This section examines the robustness of the evidence that momentum profits are only significant during optimistic investor sentiment periods. Throughout this section we continue to analyze the six-month formation and holding period strategy ($J=6$, $K=6$), and define sentiment as in Table 2.

4.3.3 Investor Sentiment, Momentum, and Market States

Cooper, Gutierrez, and Hameed (2004) propose that investors' behavioral biases will be accentuated after market gains and test whether the momentum profits are related to past market returns. They identify UP and DOWN market states using the returns of the market for a 36-month period prior to the beginning of the strategy's holding period. If this return is positive (negative), they classify the market state as UP (DOWN). Then, they compute momentum profits after UP and DOWN markets. Their results indicate that momentum profits are significant only after UP markets. This leads the authors to conclude that positive market returns amplify behavioral biases, which ultimately lead to momentum.

Market returns can, of course, be related to investor sentiment [Otoo (1999)], because, for example, as market returns increase, investors may potentially become more optimistic. However, the relationship may not be exact for two reasons. First, some investors may hold

contrarian expectations.⁴⁵ These investors may become pessimistic when they perceive that the market has climbed too high. Second, our measure of sentiment is a broad survey on aspects other than financial markets, and is likely to be affected by factors over and beyond market returns. Indeed, for our entire sample period, we find that the correlation of the time series of lagged 36-month market returns and the average residual sentiment for the past three months is 0.24. This confirms that the relationship between market returns and investor sentiment is less than perfect.

Nevertheless, a correlation of 0.24 is significant and merits investigation. Therefore, we also classify each formation period as belonging to an UP or DOWN market independently of investor sentiment as in Cooper, Gutierrez, and Hameed (2004). We calculate the return of the value-weighted index including dividends for the 36-month (Panels A1 and B1), 24-month (Panels A2 and B2) and 12-month (Panels A3 and C3) period prior to the beginning of the strategy's holding period. If this return is positive (negative), we classify the market state as UP or (DOWN). We then derive momentum profits for optimistic and pessimistic periods during UP (Panel A) and DOWN (Panel B) markets.

These results are reported in Table 4.3. From Panels A1 and B1 it can be seen that of the 500 holding period months in the sample, 436 (87.2%) occur in UP markets and only 64 (12.8%) in DOWN markets. Interestingly, in UP market states, we find considerable variation in investor sentiment, as 69 periods (or 16%) are classified as pessimistic, 254 (or 58%) as mild and 113 (or 26%) as optimistic. This provides support to the notion that market run-ups do not completely overlap with investor optimism.

Momentum strategies in DOWN markets, as shown in Panel B of Table 4.3, generally produce insignificant momentum profits, regardless of investor sentiment. However, the number of observations in each sentiment group is very small and, therefore, these results do

⁴⁵ See Grinblatt and Keloharju (2000), Chordia, Roll, and Subrahmanyam (2002), and Goetzmann and Massa (2002) for evidence on contrarian investors.

not allow meaningful interpretation. Momentum profits in UP markets, as Panel A of Table 4.3 shows, vary with investor sentiment. Specifically, in Panel A1 (36-month market return), momentum profits in optimistic periods are highly significant at a monthly average of 2.12%. This average decreases to 1.55% in mild sentiment months, and to an insignificant 0.87 (t-value = 1.51) in pessimistic months. Similar, albeit stronger, results are reported in Panels A2 (24-month market) and A3 (12-month market). These results are consistent with our previous findings, which show that momentum profits are significantly larger when investor sentiment is optimistic

For robustness we also derive the results reported in Table 4.2 using a sentiment measure that is orthogonal to the macroeconomic variables and lagged market returns. These results are shown in Table 4.4, and generally reveal the same patterns. This further demonstrates that the sentiment effect we demonstrate is not an alternative representation of the Cooper et al (2004) findings.

In Table 4.5 we report regression results. Panel A presents estimates based on the regression model of Cooper, Gutierrez, and Hameed (2004) (Table V, p. 1361), augmented with investor sentiment. Specifically, we estimate the following model (omitting time subscripts):

$$Profits = b_0 + b_1 Sentiment + b_2 Market + b_3 Market^2 + u \quad (1)$$

The variable *Profits* is the time series of average monthly momentum profits. Because we are conducting overlapping strategies, each observation of momentum profits corresponds to K formation periods, and thus K observations for investor sentiment. In the regression *Sentiment* is defined as the average sentiment calculated from these K observations. *Market* is the lagged market return of the value weighted index including

dividends during the 36, 24 and 12- month periods prior to the beginning of the strategy's holding period. $Market^2$ is the square of the market return.

The regression results in Panel A of Table 4.5 show that momentum profits increase with the market return, but decrease with the squared market term, indicating a nonlinear relationship, and confirming the results of Cooper, Gutierrez, and Hameed (2004).⁴⁶ Our results also show that the coefficient of *Sentiment* is positive and significant across all market return specifications (36, 24 and 12-months). Specifically, in Panel A, when we use a 36-month lagged market return, the coefficient on *Sentiment* is equal to 0.0002 with a *t*-value of 2.36. Whereas the magnitude of the coefficient is similar, its *t*-value increases to 3.02 when the 24-month lagged market return is used and to 3.05 when the 12-month lagged market return is used in the regressions. Likewise, the coefficient and *t*-value of the *Market* return decreases from 0.1143 with *t*-value 3.61 for the 36-month return, to 0.0666 with *t*-value 1.78 for the 12-month return. Interestingly, while the results in Panel A display that *Sentiment* predicts momentum profits independently of market returns they also show that it is a stronger predictor when the market return is calculated over shorter periods. This is consistent with the results reported in Table 4.4.

In Panels B of Table 4.5 we report results by excluding momentum profits in DOWN markets. We choose to leave out these observations because they are associated with extremely adverse market conditions, characterized by reductions in liquidity [Chordia, Roll and Subrahmanyam (2001)], and increases in volatility [Bekaert and Wu (2000)]. During such times an aggressive investment style such as momentum cannot be easily implemented. Thus, by directing our attention to UP market states, we examine the relationship between momentum profits, market returns and investor sentiment during “normal” market conditions when momentum investing can be implemented more easily.

⁴⁶ In unreported analysis (available on request) we run a regression identical to that of Cooper, Gutierrez, and Hameed (2004) (without the sentiment variables) and find results similar to theirs.

The results in Panel B demonstrate that the *only* significant variable in all three regression specifications of *Market* return (36, 24 and 12-month return) is *Sentiment*. This suggests that investor sentiment has a distinct and positive association with momentum profits during UP market states in that momentum profits in UP market conditions are related to *Sentiment* much more strongly than to past market returns.

In Panel C, we report the results of a horse race regression between optimistic sentiment, pessimistic sentiment and market returns. These results confirm those in Panel A. As before, *Sentiment* is positive and significant in all three specifications of the market return. In addition, the results show that the effect of *Market* return on momentum decreases when market returns are calculated over shorter time periods. For the 36-month period in Panel C the coefficient on *Market* return is 0.04 with t-value 2.15 and decreases to an insignificant 0.052 with t-value 1.55 for the 12-month period.

Overall the results in Tables 4.3, 4.4 and 4.5 show that the investor sentiment effect, reported in Table 4.2, is not a manifestation of the UP market effect documented by Cooper, Gutierrez, and Hameed (2004). Our findings instead suggest that investor sentiment captures a significant variation in momentum profits even after controlling for the state of the capital market.

4.3.4 Investor Sentiment, Momentum, and Trading Volume

Lee and Swaminathan (2000), argue that trading volume is related to momentum profits and show that high volume portfolios generate higher momentum returns. Trading volume has also been linked to investors' behavioral biases [Odean (1999); Statman, Thorley, and Vorkink (2006)]. In light of this evidence, we examine how sentiment affects the momentum profits generated by the different volume portfolios.

As before, we use the same methodology to construct momentum portfolios.

However, during the formation period, in addition to sorting stocks on their past returns, we also rank them on their average monthly turnover (trading volume/shares outstanding). Consistent with Lee and Swaminathan (2001), we form ten momentum groups and three trading volume groups and derive momentum returns for each combination, separately for optimistic and pessimistic formation periods.

Table 4.6 presents momentum profits for each trading volume portfolio conditional on investor sentiment. The sentiment effect is documented in all volume portfolios as momentum profits increase with sentiment. While in optimistic periods, momentum profits are highly significant for all volume portfolios, they tend to be less significant or insignificant in pessimistic periods. Interestingly, in Panel A, we observe that the high volume portfolio generates significant momentum even in pessimistic periods (0.98, t-value 2.02), approximately halved relative to optimistic periods. This result is not so surprising in light of the evidence that larger overconfidence translates to higher trading volume [Odean (1999), Grinblatt and Keloharju (2009)]. Therefore, since the high volume portfolio is indicative of larger overconfidence, it should generate momentum as suggested by the theory of Daniel et al (1998). Or, in other words, the momentum profits generated by the high volume portfolio should be less sensitive to variations in overconfidence induced by miscalibrated signals due to investor sentiment.

Mid-volume (Panel B) and low volume portfolios (Panel C), however, reveal a strong sentiment effect. In Panel B (C) for the mid (low) volume portfolio, the momentum profits decline from an average monthly return of 2.37% (2.30%) in optimistic periods to an insignificant average of 0.13% (0.04%) in pessimistic periods. In sum, while the evidence appears to be consistent with the results of Lee and Swaminathan (2000) in high volume portfolios, our results continue to highlight that momentum gains are more pronounced in optimistic states for both mild and low volume portfolios.

4.3.5 Is it a Size Effect?

A large literature suggests that return predictability is stronger for smaller companies, which are held mostly by individual investors [Nagel (2005)], and entail higher arbitrage costs [D'Avolio (2002), Jegadeesh and Titman (1993, 2001)], show that momentum strategies are more profitable amongst smaller companies. In this section, we explore whether our previous results, reported in Table 4. 2, depend on the size of the company.

We rank stocks at the end of the formation period according to firm size, and apply our momentum strategy separately to the 50% smallest and 50% largest companies.⁴⁷ These results are reported in Table 4.7 and show that sentiment affects momentum for both small and large stocks. For small stocks (Panel A), we observe that momentum profits in optimistic periods decline from a monthly average of 2.14% to 1.72% in mild sentiment periods and to an insignificant 0.46% in pessimistic periods. The corresponding figures for large companies are 0.86%, 0.81%, and 0.26%.

Our evidence that momentum is generally larger for smaller companies confirms the findings of Jegadeesh and Titman (1993, 2001). Further, the evidence that the effect of sentiment is much more dramatic in smaller companies (an average monthly return differential of 1.68% vis-à-vis 0.60% for large companies) in Panel A) supports the argument of Baker and Wurgler (2006) that the effects of investor sentiment tend to be more pronounced in the smaller companies that are harder to value and hence more prone to subjective evaluations. This momentum differential between large and small firms appears also to be consistent with the prediction of the gradual-information-diffusion theory of Hong and Stein (1999) that momentum strategies are more profitable in small firms because the information about small firms gets out more slowly and investors face fixed costs of information acquisition. Overall, the sentiment effect documented in Table 4.2 is robust to

⁴⁷ The size breakpoints are from Ken French's data library (mba.tuck.dartmouth.edu/pages/faculty/ken.french/).

firm size, as momentum profits are insignificant for both small and large companies in periods of pessimistic investor sentiment.

4.3.6 Is it Risk?

While the evidence so far suggests that conditioning on investor sentiment has a dramatic impact on the profits of momentum strategies, we cannot rule out the possibility that the higher (lower) returns of the winner (loser) portfolio during periods of optimism load more (less) strongly on economically meaningful risk factors. We address this issue by estimating CAPM, Fama and French (FF), and Conditional CAPM (CCAPM)-adjusted momentum returns across different investor sentiment states.

Following the method in Cooper, Gutierrez, and Hameed (2004), we perform the risk adjustment by forming a time series of raw momentum returns corresponding to each event month of the holding period. Specifically, to form CAPM- and FF-risk adjusted profits, for each holding period month, portfolio returns are regressed on the appropriate factors and a constant. In this manner, we obtain estimated factor loadings for each portfolio and holding period month, which we use to derive risk-adjusted profits as follows:

$$r_{kt}^{adj} = r_{kt} - \sum_i \beta_{ik} f_{it}, \quad (2)$$

where r_{kt} represents the raw returns of each momentum portfolio for the strategy in the holding period month K , in calendar month t , f_{it} is the realization of factor i in calendar month t , and β_{ik} is the estimated factor loading in month K on f_{it} . We use the excess return of the value-weighted market index, R_m , over the one-month Treasury-bill return, R_f as the market portfolio in the CAPM, and, additionally, the return differential between small and big companies (SMB), and high and low book-to-market companies (HML), for the FF risk

adjustment.⁴⁸

For the CCAPM we allow the covariance between the returns of momentum portfolios with the excess market return to vary with investor sentiment. Particularly we estimate risk adjusted returns using the following model:

$$r_{kt}^{adj} = r_{kt} - (\beta_{ik} - \beta_{ik}^{sent} * Sentiment_{t-j})(R_m - R_f), \quad (3)$$

where r_{kt} represents the raw returns of each momentum portfolio for the strategy in the holding period month K , in calendar month t , β_{ik} is the estimated factor loading in month K on the excess market return and β_{ik}^{sent} is the factor loading in month K on the interaction between the excess market return and investor sentiment during the formation period.⁴⁹ The time-varying betas argument predicts that the covariance between momentum profits and excess market returns increases when sentiment is optimistic; therefore returns increase accordingly to compensate for the increase in the co-variation between momentum portfolios and the excess market return.⁵⁰

Table 4.8 shows the CAPM, FF and CCAPM-adjusted momentum profits. The pattern of momentum profits, reported in Table 2, remains robust to these risk adjustments. Momentum profits are highly significant, at a monthly average of 2.03% (CAPM), 2.08% (FF) and 2.03% (CCAPM), respectively, when the strategy is implemented in optimistic investor sentiment periods. However, in pessimistic periods momentum profits drop remarkably to a monthly average return of 0.48% (CAPM), 0.96% (FF) and 0.47% (CCAPM), respectively. Note that the CAPM and the CCAPM-adjusted returns are virtually

⁴⁸ We thank Eugene Fama and Kenneth French for providing the data on Fama and French (1993) factors on the WRDS website.

⁴⁹ Because we perform overlapping strategies for each portfolio return observation we have CB residuals from K formation periods. In Equation (4), *Sentiment* is the average sentiment from these K formation periods.

⁵⁰ We do not perform a conditional FF specification because the SMB and HML factors may be related to sentiment in a manner that is consistent with a behavioral story. Therefore, allowing factor loadings between momentum returns and the HML and SMB portfolios to vary according to investor sentiment will produce inconclusive results.

indistinguishable, suggesting that beta does not depend on investor sentiment. This result is in line with the findings of Baker and Wurgler (2006).

Overall, it is reasonable to conclude that rational risk premia, at least in the context of the CAPM and the Fama and French (1993) models, are not able to explain the superior performance of momentum strategy in periods of optimistic investor sentiment.

4.3.7 Is the CB Sentiment Index Forecasting Future Macroeconomic Conditions?

Although the sentiment effect on momentum profits cannot be explained by traditional asset pricing theories it may still be related to some form of rational pricing if our sentiment measure reflects future macroeconomic activity, which determines holding period returns. In this section we construct our sentiment measure by controlling for proxies for future macroeconomic conditions. Specifically, we replicate the analysis in Panel A of Table 4.2 by orthogonalizing the CB index on current macroeconomic indicators (described in Section 1), their one-quarter ahead values,⁵¹ and the closing value of the VIX at the end of the month in which sentiment is measured.⁵² The inclusion of the one-quarter ahead macroeconomic variables allows us to control for future macroeconomic conditions. Since the VIX is obtained primarily from the activity of sophisticated investors [Lakonishok, Pearson and Poteshman (2007)] who participate in the options market, it is included in the regression as a proxy for macroeconomic expectations.⁵³ All of the other calculations remain

⁵¹ Only the current value of the NBER recession dummy is included.

⁵² VIX stands for Market Volatility Index, and is based on index option prices. For a general discussion on the VIX see Whaley (2008).

⁵³ The level of the VIX may also be related to sentiment, as shown by Kaplanski and Levi (2009). However, since options are primarily traded by sophisticated investors it is reasonable to assume that the sentiment displayed by the VIX is not as dramatic as that displayed by the CB index. Therefore, at least, some of the variability in the VIX is related to expectations about future macroeconomic conditions.

unchanged, with the difference that the sample period of this test is shorter as the VIX is available from January 1985.⁵⁴

Table 4.9 reports momentum returns for $J=6, K=3$ (Panel A), $J=K=6$ (Panel B) and $J=6, K=12$ (Panel C) strategies. The new results remain essentially unchanged even when we orthogonalize the CB index with respect to future macroeconomic conditions and VIX. For example, as seen in Panel B, momentum profits in optimistic sentiment states are highly significant at a monthly average of 2.42%. They decline to 1.51% in mild sentiment periods and to an insignificant 0.27% in pessimistic periods. Similar results obtain in Panels A and C. These findings suggest that the sentiment measure used in our analysis is not forecasting future macroeconomic activity, which, in turn, explains momentum profits.

4.3.8 An Alternative Sentiment Index

In this section, we examine the sensitivity of our results to an alternative index for investor sentiment using the monthly measure constructed by Baker and Wurgler (2006, 2007).⁵⁵ These authors suggest that investor sentiment can be captured from various market-based variables that relate to investors' propensity to purchase stocks. They construct a sentiment time series using six sentiment-revealing variables: trading volume (measured as total NYSE turnover),⁵⁶ dividend premium, closed-end fund discount, number and first day returns in IPOs, and the equity share in new issues. Because these variables are partly related to economic fundamentals, they regress each of these sentiment proxies against growth in industrial production, real growth in durable consumption, non-durable consumption, services consumption, growth in employment, and an NBER recession indicator, and use the residuals

⁵⁴ We use both the ways to construct the VIX described in Whaley (2008, footnote 9 pp. 4) and get identical results. The reported results are based on the first method of Whaley (1988).

⁵⁵ This index is available from Jeffrey Wurgler's website (<http://pages.stern.nyu.edu/~jwurgler/>).

⁵⁶ To remove the time trend from the turnover, Baker and Wurgler (2006, 2007) use log turnover minus a five year moving average.

from this regression as the sentiment proxies. The overall sentiment index is the first principal component of the six sentiment proxies. For more detail on the construction of the index, see Baker and Wurgler (2006, 2007). This time series is available on a monthly basis from October 1965 to December 2007.

Table 4.10, reports Panel A, Table 2-equivalent momentum results for optimistic and pessimistic periods, using the Baker and Wurgler sentiment measure. Aside from the replacement of the CB index with the Baker and Wurgler measure, all calculations remain the same as those in Table 2, with the small difference that in this table we combine the Mild sentiment category with the Optimistic sentiment category. We do this because Mild and Optimistic sentiment states yield very similar momentum profits, thus by pooling them together we gain statistical power without losing important information.⁵⁷

Consistent with our earlier baseline findings, the new evidence confirms the difference in momentum profits between optimistic and pessimistic investor states even when we use an alternative investor sentiment index. Specifically, in Panel B ($J,K=6$), these results show that momentum profits in optimistic periods are equal to an average monthly return of 1.59%, whereas in pessimistic periods they drop to an insignificant 0.30%. Similar results are seen in Panel A for the $J=6,K=3$ (Panel B: $J=6,K=12$) momentum strategy where the corresponding momentum profits in optimistic periods are 1.71% (0.97%), reducing to an insignificant 0.23% (-0.01%) in pessimistic periods.⁵⁸ These findings corroborate our previous results and show that they are not driven by the choice of sentiment index.

⁵⁷ When momentum profits are partitioned to three sentiment categories, and $J,K=6$, the optimistic sentiment category yields average momentum profits equal 1.40% and the mild sentiment category 1.67%. The corresponding figures for Panels A ($K=3$) and C ($K=12$) are 1.54% and 1.80% and 0.70% and 1.05%.

⁵⁸ We obtain very results when we use 20%-20% sentiment states.

4.3.9 Alternative Lags for Optimistic and Pessimistic Sentiment

In our analysis so far, we classify each formation period as pessimistic or optimistic using a rolling average of the residual sentiment level during a three-month window prior to the beginning of the holding period. In this section, we examine the sensitivity of our results to average sentiment calculated as the average of the two and four months prior to the end of the formation period.⁵⁹

As shown in Panels A (2 sentiment lags) and B (4 sentiment lags) of Table 4.11, our main results hold for this alternative sentiment specification. Momentum strategies in optimistic periods consistently yield significant average monthly profits of 1.95% (Panel A) and 2.10% (Panel B). These profits, however, decline substantially in pessimistic periods, equaling 0.46% (Panel A), and 0.62% (Panel B). Overall, these results confirm that our baseline findings are robust to a different definition of investor sentiment.

4.4 Momentum Profits, Investor Sentiment and Long-Run Returns

A central prediction of behavioral theories such as Daniel, Hirshleifer, and Subrahmanyam (1998) is that momentum profits reflect unrealistic expectations, and thus revert in the long run. Since in the previous section we documented that momentum profits are *only* significant when investors are optimistic, we would expect these profits to reverse over longer horizons. In this section, we examine the pattern of momentum profits in event time, six years (Panel A), 5 years (Panel B), 4 years (Panel C) and 3 years (Panel D) after portfolio formation.

Table 4.12 presents the results. As expected, momentum profits revert *only* after optimistic periods, regardless of whether returns are risk-adjusted. For portfolios constructed

⁵⁹ The calculation of these rolling averages uses equal weights.

in optimistic formation periods using raw returns, this reversal, as shown in Panel A1, is equal to an average monthly return of -0.56%. The corresponding figure for CAPM (FF)-adjusted returns is -0.56% (-0.45%). However, for portfolios formed in pessimistic periods, as expected, there is no reversal. The momentum returns are equal to -0.02% and 0.00% for raw and CAPM adjusted returns respectively, and 0.11% for Fama-French adjusted returns. Similar findings are shown in Panels B,C and D, confirming that our conclusion is not dependent on the holding horizon.

An noteworthy result that surfaces in Table 4.12 is that while mild sentiment portfolios experience significant reversals when we focus on raw (-0.17%) or CAPM (-0.18%) adjusted returns, these momentum profits become insignificant when they are adjusted using the Fama-French 3 factor model (-0.02%). This result suggests that the reversal in the mild sentiment portfolio is related to the HML and/or SMB factor, and thus is not abnormal in the context of the 3-factor model.⁶⁰ This finding appears to be in line with Fama and French (1996), who show that the 3-factor model can price the long-run reversals documented by De bondt and Thaler (1985). However, the FF 3-factor model cannot explain the economically and statistically significant reversals in the optimistic sentiment momentum portfolios.

Another striking pattern in Table 11 is that, regardless of the risk adjustment technique, the monotonic pattern in the short-run returns of the 10 *optimistic* momentum portfolios (i.e., returns increasing from the loser to the winner portfolio) *completely* reverts in the long-run as returns are decreasing from the loser to the winner portfolio. A similar, albeit less pronounced, pattern is seen in mild sentiment periods. On the contrary, however, the long-run returns of the pessimistic momentum portfolios are virtually indistinguishable.

⁶⁰ This does not necessarily mean it is rational in the sense that its *strictly* driven by covariance with state variables, as there is evidence suggesting that returns are affected by characteristics (Daniel and Titman 1997), and/or mispricing (Daniel, Hirshleifer and Subrahmanyam 2001), rather than factor loadings.

This finding reveals an intimate link between short-run momentum and long-run reversal. The trading actions of optimistic investors lead to short-run momentum by pushing prices away from fundamental values. As these expectations fail to materialize, investor sentiment subsides with momentum profits fading away, and stock prices reverting to fundamental values in the long run. This finding provides support to the behavioral model of Daniel, Hirshleifer, and Subrahmanyam (1998), which predicts that short-run momentum and long-run stock price reversal commonly arise from investors' behavioral biases.

These results provide an important link between short-run price momentum and long-run reversal. Cooper, Gutierrez, and Hameed (2004) and Lee and Swaminathan (2000) document such links. The former authors show that momentum profits revert after UP markets, where short-run momentum is significant. However, they also find that momentum profits revert after DOWN markets, and the difference in the reversals between UP and DOWN markets is not significant. Lee and Swaminathan (2000) find that trading volume also predicts reversals, albeit differently for winners and losers.⁶¹ Our study corroborates this evidence by showing that investor sentiment predicts a significant difference in the long-run performance of momentum portfolios.

4.5 Further discussion and concluding remarks.

Does investor sentiment affect financial asset prices? This issue is enduring and has taken on renewed significance in the context of dramatic rises and falls in the stock market during this decade. For this reason both academics and practitioners are taking an interest in how investors' psychology can affect financial markets.

Price momentum is perhaps the most puzzling pattern in the data. Our analysis

⁶¹ They find that momentum portfolios comprised of high volume winners and low volume losers exhibit reversals, whereas the opposite classifications result to continuations.

provides support to the notion that investors' sentiment does affect the stock market, and thus should be accounted for when attempting to forecast stock returns. In addition, our results highlight the role of limits to arbitrage, in the persistence of mispricings.

Specifically, we show that momentum is only significant when investors are optimistic. These are periods where stocks are generally overpriced. Traditional arbitrage arguments, would suggest that rational investors would spot the mispricing and short sell the overpriced stocks. However, because arbitrage costs (in the form of fees and noise trader risk) are high, rational investors cannot do this effectively, allowing the momentum mispricing to persist. However, in the long run prices slowly gravitate to fundamental values, therefore causing reversals.

As noted by Shleifer and Vishny (1997) patterns in the data that seem like profit opportunities (like momentum) persist because of high arbitrage costs. In other words, arbitragers do spot them, but because it is very costly to arbitrage these portfolios, it is not forceful enough to completely eliminate the mispricing. Our results, coupled with those of Hvidkjaer (2006) who finds that momentum is related to individual traders (Hvidkjaer 2006) and Jegadeesh and Titman (1993, 2001) that momentum is stronger amongst smaller stocks which are more difficult to arbitrage, support this interpretation and highlight the intimate interplay between investors' psychological biases and the limits to arbitrage as the drivers of price predictability.

Figure 1
Investor sentiment (CB) from 1966-2008

This figure plots three series. The first is the raw data of consumer confidence provided by the Conference Board. The second series is the 3-month rolling average of the residual from regressing the CB index series on the following set of macroeconomic variables: growth in industrial production, real growth in durable consumption, non-durable consumption, services consumption, growth in employment, and an NBER recession indicator.




Table 4.1 Descriptive Statistics

Panel A presents descriptive statistics for the raw time series of consumer confidence, as compiled by Conference Board. Panel B presents the 3 month rolling average of the component of investor sentiment that is orthogonal to macroeconomic conditions. To derive this component we regress raw sentiment on growth in industrial production, real growth in durable, non-durable, and services consumption, growth in employment, and an NBER recession indicator, and then use the residuals from this regression to calculate the 3-month rolling average. The sample period is April 1967 to December 2008.

Panel A: CB consumer confidence

Mean	σ	Q1	Median	Q3	Minimum	Maximum	N
97.40	23.06	82.59	98.00	110.60	38.62	144.71	501

Panel B: CB consumer confidence orthogonal to macroeconomic variables

-0.06	19.96	-13.13	-1.32	11.91	-52.89	44.13	501
-------	-------	--------	-------	-------	--------	-------	-----

Table 4.2
Momentum Profits Conditional on Investor Sentiment

This table presents average monthly returns in percentages for price momentum strategies involving all NYSE/AMEX stocks for the time period April 1967 until December 2008. At the beginning of each month all stocks are ranked based on their cumulative returns over the previous J months. Portfolio 1 includes the loser stocks and portfolio 10 the winner stocks. The winner stocks are bought and the loser stocks sold, and this position is held for K months. Monthly holding period returns come from overlapping strategies and are computed as an equal-weighted average of returns from strategies initiated at the beginning of this month, and the previous $K-1$ months. We allow one month between the end of the formation period and the beginning of the holding period, and delete all stocks that are priced less than one \$1 at the beginning of the holding period. Sentiment is measured using the time series of consumer confidence sentiment index constructed by Conference Board. We regress this series on growth in industrial production, real growth in durable, non-durable, and services consumption, growth in employment, and an NBER recession indicator, and use the residuals from this regression as the sentiment proxy. In order to identify whether a particular formation period was optimistic or pessimistic, in each month t we calculate the average sentiment level for the previous 3 months, using a weight of 3/6 for month t , a weight of 2/6 for month $t-1$ and a weight of 1/6 for month $t-2$. In Panel A (B) (C) the top 30% (20%) (40%) observations of this rolling average time series are the high sentiment periods, and the bottom 30% (20%) (40%) the low sentiment periods. To identify each holding period month as optimistic and pessimistic, we calculate how many of the K formation periods were of high and low sentiment. If all K formation periods were classified as optimistic (pessimistic) the holding period month is classified as optimistic (pessimistic). To test whether momentum profits in each sentiment state respectively are equal to zero, we regress the time series of average monthly momentum profits on Optimistic, Pessimistic and Mild sentiment dummies, with no intercept. To test if mean profits in Optimistic sentiment periods are different from profits in Pessimistic sentiment periods we regress average monthly momentum profits on a Mild sentiment dummy variable and an Optimistic sentiment dummy variable with a constant. The t -statistics of the significance of momentum profits and the difference between profits derived after optimistic and pessimistic periods are calculated using Newey-West standard errors, where the lag is set to $K-1$.

		Momentum Portfolio										
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell [t-stat.]
Panel A: 30%-30% Sentiment States												
Panel A1: J=6, K=3												
Optimistic	(n=133)	-0.85	-0.05	0.36	0.53	0.68	0.78	0.96	0.90	1.09	1.42	2.27 [5.06]
Mild	(n=246)	-0.24	0.53	0.81	0.91	0.97	1.01	1.04	1.04	1.14	1.20	1.44 [4.67]
Pessimistic	(n=121)	1.86	2.07	2.21	2.14	2.08	1.97	1.91	1.84	1.89	2.15	0.29 [0.56]
											Opt.-Pes.	1.98 [2.90]
Panel A2: J,K=6												
Optimistic	(n=121)	-0.48	0.17	0.52	0.67	0.86	0.95	1.10	1.10	1.26	1.51	2.00 [5.66]
Mild	(n=286)	-0.28	0.34	0.63	0.77	0.86	0.91	0.99	1.05	1.12	1.18	1.46 [5.66]
Pessimistic	(n=93)	2.12	2.41	2.45	2.31	2.24	2.21	2.11	2.13	2.21	2.45	0.34 [0.77]
											Opt.-Pes.	1.66 [2.96]
Panel A3 :J=6,K=12												
Optimistic	(n=109)	-0.27	0.20	0.49	0.62	0.78	0.86	0.94	0.96	1.00	1.00	1.27 [3.50]
Mild	(n=337)	0.38	0.75	0.95	1.02	1.10	1.12	1.18	1.22	1.25	1.23	0.85 [4.14]
Pessimistic	(n=54)	2.05	2.25	2.21	2.12	2.11	2.05	2.00	2.00	2.06	2.14	0.09 [0.22]
											Opt.-Pes.	1.18 [2.20]

Table 4.2 , continued

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
Panel B: 20%-20% Sentiment States													
Panel B1: J=6, K=3													
Optimistic	(n=92)	-0.67	-0.09	0.30	0.47	0.60	0.71	1.01	0.87	1.10	1.55	2.22	[3.83]
Mild	(n=327)	-0.24	0.56	0.86	0.98	1.04	1.06	1.09	1.08	1.14	1.22	1.46	[5.10]
Pessimistic	(n=81)	2.21	2.31	2.36	2.20	2.14	2.03	1.90	1.91	2.07	2.32	0.11	[0.25]
											Opt.-Pes	2.11	[2.87]
Panel B2: J,K=6													
Optimistic	(n=84)	-0.43	-0.04	0.25	0.40	0.58	0.67	0.91	0.85	1.08	1.44	1.87	[4.17]
Mild	(n=353)	-0.06	0.60	0.90	1.03	1.12	1.16	1.20	1.27	1.30	1.38	1.44	[6.24]
Pessimistic	(n=63)	1.80	2.11	2.07	1.89	1.84	1.83	1.77	1.77	2.02	2.20	0.40	[1.04]
											Opt.-Pes.	1.47	[2.48]
Panel B3 :J=6, K=12													
Optimistic	(n=100)	0.04	0.42	0.66	0.76	0.91	0.95	1.06	1.09	1.11	1.18	1.14	[3.11]
Mild	(n=301)	0.23	0.65	0.89	0.99	1.07	1.11	1.19	1.24	1.26	1.24	1.01	[4.62]
Pessimistic	(n=99)	1.40	1.61	1.62	1.55	1.56	1.53	1.47	1.46	1.51	1.51	0.11	[0.41]
											Opt.-Pes.	1.03	[2.28]

Table 4.2 , continued

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
		Panel C: 40%-40% Sentiment State											
Panel C1:J=6,K=3													
Optimistic	(n=172)	-1.51	-0.53	-0.07	0.12	0.29	0.39	0.57	0.52	0.65	0.87	2.38	[6.09]
Mild	(n=170)	0.24	0.91	1.16	1.27	1.31	1.33	1.35	1.35	1.42	1.48	1.24	[2.97]
Pessimistic	(n=158)	1.63	1.88	2.00	1.94	1.89	1.81	1.75	1.71	1.83	2.08	0.45	[1.05]
											Opt.-Pes	1.93	[3.33]
Panel C2: J,K=6													
Optimistic	(n=149)	-0.99	-0.23	0.17	0.36	0.58	0.68	0.85	0.90	1.05	1.26	2.25	[6.83]
Mild	(n=227)	0.05	0.65	0.93	1.05	1.08	1.11	1.15	1.18	1.24	1.26	1.21	[4.12]
Pessimistic	(n=124)	1.56	1.82	1.89	1.82	1.83	1.85	1.80	1.85	1.93	2.21	0.65	[1.83]
											Opt.-Pes.	1.60	[3.29]
Panel C3 :J=6,K=12													
Optimistic	(n=124)	-0.65	-0.08	0.24	0.38	0.57	0.66	0.76	0.79	0.82	0.80	1.45	[3.91]
Mild	(n=304)	0.50	0.83	1.04	1.11	1.18	1.20	1.26	1.30	1.32	1.31	0.81	[3.71]
Pessimistic	(n=72)	1.96	2.13	2.06	1.95	1.93	1.89	1.84	1.84	1.89	2.00	0.04	[0.14]
											Opt.-Pes.	1.41	[2.87]

Table 4.3
Momentum Profits Conditional on Different Market States and Investor Sentiment

This table presents average monthly returns in percentages for price momentum strategies involving all NYSE/AMEX stocks for the time period April 1967 until December 2008. Panel A shows momentum strategies implemented in UP markets, whereas Panel B momentum strategies implemented after DOWN markets. The state of the market is the return of the value weighted market index including dividends 36 (Panel A1,B1), 24 (A2,B2) and 12 (A3,B3) months prior to beginning of the holding period, as measured by Cooper et al (2004). We allow one month between the end of the formation period and the holding period, and delete all stocks that are priced less than one \$1 at the beginning of the holding period. Sentiment is defined as in table 2. To test whether momentum profits in each sentiment state respectively are equal to zero, we regress the time series of average monthly momentum profits on Optimistic, Pessimistic and Mild sentiment dummies, with no intercept. To test if mean profits in Optimistic sentiment periods are different from profits in Pessimistic sentiment periods we regress average monthly momentum profits on a Mild sentiment dummy variable and an Optimistic sentiment dummy variable with a constant. In this table $J, K=6$.

Momentum Portfolio													
<i>Panel A: UP markets</i>													
Panel A1: 36-month market		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
Optimistic	(n=113)	-0.60	0.11	0.48	0.63	0.82	0.90	1.05	1.08	1.24	1.52	2.12	[4.73]
Mild	(n=254)	-0.27	0.39	0.70	0.83	0.95	1.00	1.06	1.14	1.22	1.28	1.55	[5.18]
Pessimistic	(n=69)	1.15	1.74	1.88	1.80	1.75	1.73	1.61	1.62	1.74	2.01	0.87	[1.51]
											Opt.-Pes.	1.25	[1.72]
Panel A2: 24-month market													
Optimistic	(n=103)	-0.75	-0.02	0.36	0.54	0.75	0.82	0.98	1.01	1.18	1.45	2.21	[4.86]
Mild	(n=252)	-0.34	0.37	0.69	0.83	0.95	1.01	1.09	1.18	1.29	1.40	1.74	[6.01]
Pessimistic	(n=76)	1.50	1.94	2.00	1.91	1.86	1.84	1.74	1.74	1.89	2.16	0.66	[1.25]
											Opt.-Pes.	1.55	[2.22]
<i>Panel A3: 12-month market</i>													
Optimistic	(n=87)	-0.50	0.30	0.71	0.85	1.08	1.15	1.34	1.42	1.62	2.03	2.53	[6.16]
Mild	(n=206)	0.06	0.69	0.93	1.00	1.09	1.12	1.17	1.24	1.36	1.49	1.43	[5.36]
Pessimistic	(n=86)	1.81	2.22	2.29	2.16	2.11	2.10	2.01	2.01	2.14	2.42	0.61	[1.49]
											Opt.-Pes.	1.92	[3.29]

Table 4.3, continued

Momentum Portfolio													
<i>Panel B: DOWN markets</i>													
Panel B1: 36-month market		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
Optimistic	(n=8)	1.20	0.95	1.00	1.31	1.43	1.57	1.80	1.39	1.53	1.46	0.26	[0.10]
Mild	(n=32)	-0.36	-0.11	0.02	0.37	0.22	0.24	0.38	0.35	0.29	0.40	0.76	[0.60]
Pessimistic	(n=24)	4.91	4.33	4.09	3.78	3.65	3.60	3.55	3.59	3.55	3.72	-1.20	[-0.82]
											Opt.-Pes.	1.46	[0.50]
Panel B2: 24-month market													
Optimistic	(n=18)	1.06	1.27	1.40	1.44	1.49	1.66	1.79	1.59	1.71	1.87	0.80	[0.45]
Mild	(n=34)	0.17	0.11	0.18	0.38	0.21	0.19	0.25	0.05	-0.19	-0.45	-0.62	[-0.48]
Pessimistic	(n=17)	4.88	4.52	4.47	4.11	3.93	3.86	3.77	3.86	3.65	3.75	-1.12	[0.61]
											Opt.-Pes.	1.92	[0.75]
<i>Panel B3: 12-month market</i>													
Optimistic	(n=34)	-0.45	-0.16	0.03	0.22	0.32	0.43	0.49	0.27	0.33	0.18	0.63	[0.47]
Mild	(n=80)	-1.17	-0.58	-0.15	0.20	0.27	0.38	0.52	0.57	0.50	0.37	1.54	[1.74]
Pessimistic	(n=7)	5.95	4.67	4.42	4.15	3.81	3.60	3.41	3.54	3.02	2.86	-3.09	[-1.03]
											Opt.-Pes.	3.72	[1.13]

Table 4.4
Momentum Profits Conditional on Investor Sentiment orthogonal to macroeconomic variables and lagged market returns

This table presents average monthly returns in percentages for price momentum strategies involving all NYSE/AMEX stocks for the time period April 1967 until December 2008. At the beginning of each month all stocks are ranked based on their cumulative returns over the previous J months. Portfolio 1 includes the loser stocks and portfolio 10 the winner stocks. The winner stocks are bought and the loser stocks sold, and this position is held for K months. Monthly holding period returns come from overlapping strategies and are computed as an equal-weighted average of returns from strategies initiated at the beginning of this month, and the previous $K-1$ months. We allow one month between the end of the formation period and the beginning of the holding period, and delete all stocks that are priced less than one \$1 at the beginning of the holding period. Sentiment is measured using the time series of consumer confidence sentiment index constructed by Conference Board. We regress this series on growth in industrial production, real growth in durable, non-durable, and services consumption, growth in employment, an NBER recession indicator, and the lagged market return for 36 months (Panel A), 24 months (Panel B) and 12 months (Panel C) ending in the month in which sentiment is measured. The residuals from this regression are the sentiment proxy. The remaining calculations are as in Table 4.2

		Momentum Portfolio												
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	t-stat.	
Panel A: 36 month market														
<i>Panel A1: J=6, K=3</i>														
Optimistic	(n=133)	-1.13	-0.31	0.09	0.31	0.47	0.57	0.76	0.68	0.87	1.16	2.29	[5.03]	
Mild	(n=252)	0.41	1.01	1.26	1.34	1.34	1.34	1.39	1.39	1.47	1.59	1.18	[3.62]	
Pessimistic	(n=115)	0.75	1.28	1.47	1.43	1.47	1.43	1.31	1.28	1.37	1.53	0.78	[1.83]	
												Opt.-Pes	1.51	[2.43]
<i>Panel A2: J, K=6</i>														
Optimistic	(n=117)	-0.47	0.14	0.48	0.65	0.84	0.93	1.07	1.07	1.24	1.48	1.95	[3.11]	
Mild	(n=296)	0.21	0.71	0.97	1.07	1.12	1.15	1.21	1.27	1.34	1.43	1.22	[4.56]	
Pessimistic	(n=87)	0.58	1.29	1.46	1.44	1.51	1.52	1.45	1.47	1.56	1.76	1.17	[3.11]	
												Opt.-Pes.	0.78	[1.49]
<i>Panel A3 : J=6, K=12</i>														
Optimistic	(n=103)	-0.18	0.19	0.45	0.57	0.74	0.82	0.90	0.92	0.96	0.98	1.16	[3.06]	
Mild	(n=346)	0.40	0.79	1.02	1.09	1.16	1.19	1.26	1.30	1.32	1.30	0.90	[4.12]	
Pessimistic	(n=51)	1.82	1.99	1.88	1.73	1.77	1.68	1.58	1.56	1.68	1.79	-0.03	[-0.07]	
												Opt.-Pes.	1.19	[2.07]

Table 4.4 Continued

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
Panel B: 24-month market													
<i>Panel B1: J=6, K=3</i>													
Optimistic	(n=135)	-1.10	-0.26	0.15	0.36	0.51	0.61	0.82	0.73	0.92	1.22	2.33	[5.18]
Mild	(n=244)	0.18	0.85	1.11	1.19	1.20	1.20	1.24	1.24	1.28	1.35	1.18	[3.59]
Pessimistic	(n=121)	1.19	1.57	1.74	1.68	1.71	1.68	1.58	1.56	1.70	1.94	0.74	[1.89]
											Opt.-Pes	1.59	[2.65]
<i>Panel B2: J, K=6</i>													
Optimistic	(n=120)	-0.65	0.03	0.37	0.56	0.74	0.83	0.98	0.98	1.14	1.36	2.01	[5.66]
Mild	(n=292)	0.14	0.64	0.92	1.02	1.09	1.13	1.20	1.26	1.33	1.42	1.28	[4.80]
Pessimistic	(n=88)	1.07	1.69	1.78	1.71	1.73	1.75	1.64	1.66	1.74	1.94	0.88	[2.54]
											Opt.-Pes.	1.13	[2.28]
<i>Panel B3 : J=6, K=12</i>													
Optimistic	(n=106)	-0.23	0.19	0.46	0.59	0.74	0.82	0.91	0.93	0.97	0.99	1.22	[3.29]
Mild	(n=345)	0.51	0.87	1.07	1.15	1.22	1.24	1.30	1.35	1.37	1.36	0.85	[3.96]
Pessimistic	(n=49)	1.24	1.54	1.53	1.39	1.41	1.36	1.30	1.26	1.33	1.39	0.15	[0.44]
											Opt.-Pes.	1.07	[2.10]

Table 4.4 Continued

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
Panel C: 12 month market													
<i>Panel C1: J=6, K=3</i>													
Optimistic	(n=133)	-0.85	-0.05	0.36	0.53	0.68	0.78	0.96	0.90	1.09	1.42	2.27	[5.06]
Mild	(n=245)	-0.28	0.49	0.77	0.88	0.94	0.98	1.01	1.01	1.11	1.16	1.45	[4.68]
Pessimistic	(n=122)	1.82	2.03	2.17	2.10	2.04	1.93	1.87	1.81	1.85	2.10	0.28	[0.55]
											Opt.-Pes	1.99	[2.93]
<i>Panel C2: J, K=6</i>													
Optimistic	(n=121)	-0.48	0.17	0.52	0.67	0.86	0.95	1.10	1.10	1.26	1.51	2.00	[5.66]
Mild	(n=285)	-0.30	0.33	0.62	0.77	0.86	0.91	0.98	1.04	1.11	1.17	1.46	[5.65]
Pessimistic	(n=94)	2.13	2.42	2.45	2.31	2.25	2.22	2.11	2.13	2.21	2.47	0.34	[0.79]
											Opt.-Pes.	1.66	[2.97]
<i>Panel C3 : J=6, K=12</i>													
Optimistic	(n=109)	-0.27	0.20	0.49	0.62	0.78	0.86	0.94	0.96	1.00	1.00	1.27	[3.50]
Mild	(n=337)	0.38	0.75	0.95	1.02	1.10	1.12	1.18	1.22	1.25	1.23	0.85	[4.14]
Pessimistic	(n=54)	2.05	2.25	2.21	2.12	2.11	2.05	2.00	2.00	2.06	2.14	0.09	[0.22]
											Opt.-Pes.	1.18	[2.20]

Table 4.5
Regressions of Momentum Profits on Market Returns and Investor Sentiment

Market is the return of the value weighted market index 36, 24 and 12 months prior to beginning of the holding period, and market return² is the square term of the market return. For each momentum profit observation (which corresponds to K formation periods due to overlapping strategies) we calculate average sentiment (defined as in Table 2) in the K formation periods. T- statistics are calculated using Newey-West standard errors, where the lag is set to $K-1$. In this table $J, K=6$.

	Parameter	36-month market return			24-month market return			12-month market return		
		Estimate	t- statistic	Adj.R ²	Estimate	t- statistic	Adj.R ²	Estimate	t- statistic	Adj.R ²
Panel A: Cooper et al regression with sentiment: Mom. profits = $b_0 + b_1 \cdot \text{Sentiment} + b_2 \cdot \text{Market} + b_3 \cdot \text{Market}^2 + u$										
Constant	b_0	0.0099	3.13	0.0291	0.0137	4.58	0.026	0.0141	5.02	0.013
Sentiment	b_1	0.0002	2.36		0.0003	3.02		0.0003	3.05	
Market return	b_2	0.1143	3.61		0.0905	2.50		0.0666	1.78	
Market return ²	b_3	-0.3804	-3.30		-0.4783	-2.51		-0.480	-1.64	
Panel B: Cooper et al regression with sentiment in UP markets: Mom. profits = $a_0 + b_1 \cdot \text{Sentiment} + b_3 \cdot \text{UPmarket ret.} + b_4 \cdot \text{UPmarket ret.}^2 + u$										
Constant	b_0	0.0027	0.31	0.005	0.0103	1.29	0.009	0.0089	1.39	0.022
Sentiment	b_1	0.0002	1.75		0.0002	2.15		0.0003	2.70	
Market return	b_2	0.1831	1.66		0.1819	1.24		0.2211	1.40	
Market return ²	b_3	-0.5259	-1.66		-0.8778	-1.49		-1.406	-1.59	
Panel C: Horse race between market returns and sentiment: Mom. profits = $b_0 + b_1 \cdot \text{Market} + b_2 \cdot \text{Sentiment} + u$										
Constant	b_0	0.0076	2.01	0.015	0.0093	2.72	0.014	0.0115	4.17	0.010
Market return	b_1	0.0476	2.15		0.0520	1.92		0.0525	1.55	
Sentiment	b_2	0.0002	2.27		0.0002	2.80		0.0003	3.07	

Table 4.6
Momentum Profits Conditional on Investor Sentiment and Trading Volume

This table presents average monthly returns in percentages for price momentum strategies involving all NYSE/AMEX stocks for the time period April 1967 until December 2008. At the end of the formation period we rank stocks in deciles based on cumulative returns in the previous J months, as well as their average monthly turnover (total volume/shares outstanding) over the same period. We then independently form 10 portfolios based on past returns and 3 portfolios based on volume, which we hold for K months using overlapping strategies. Sentiment is defined as in Table 2. In this table we use 30%-30% cut-off points for optimistic and pessimistic sentiment. To test whether momentum profits in each sentiment state respectively are equal to zero, we regress the time series of average monthly momentum profits on Optimistic, Pessimistic and Mild sentiment dummies, with no intercept. To test if mean profits in Optimistic sentiment periods are different from profits in Pessimistic sentiment periods we regress average monthly momentum profits on a Mild sentiment dummy variable and an Optimistic sentiment dummy variable with a constant. The t -statistics of the significance of momentum profits after optimistic and pessimistic sentiment in the volume portfolios, and the difference between profits derived after optimistic and pessimistic periods are calculated using Newey-West standard errors, where the lag is set to $K-1$. In this table $J, K=6$.

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
Panel A: High Volume													
Optimistic	(n=121)	-0.50	0.12	0.28	0.29	0.50	0.58	0.64	0.65	0.91	1.32	1.82	[4.11]
Mild	(n=286)	-0.52	0.22	0.47	0.61	0.71	0.77	0.88	0.88	0.99	1.05	1.57	[5.13]
Pessimistic	(n=93)	1.27	1.87	2.00	1.84	1.81	1.87	1.78	1.86	1.94	2.25	0.98	[2.02]
											Opt.-Pes.	0.84	[1.27]
Panel B: Mid-Volume													
Optimistic	(n=121)	-0.49	0.49	0.88	0.85	1.00	1.06	1.15	1.18	1.44	1.87	2.37	[6.99]
Mild	(n=286)	-0.32	0.31	0.63	0.78	0.86	0.93	0.95	1.07	1.12	1.31	1.62	[6.16]
Pessimistic	(n=93)	2.44	2.42	2.38	2.24	2.24	2.12	2.07	2.08	2.15	2.58	0.13	[0.32]
											Opt.-Pes.	2.24	[4.16]
Panel C: Low Volume													
Optimistic	(n=121)	0.36	0.24	0.54	0.79	1.02	1.04	1.18	1.32	1.44	1.94	2.30	[6.37]
Mild	(n=286)	0.18	0.57	0.80	0.92	0.97	1.00	1.09	1.15	1.26	1.32	1.14	[4.60]
Pessimistic	(n=93)	2.85	2.76	2.75	2.60	2.47	2.46	2.37	2.53	2.70	2.90	0.04	[0.10]
											Opt.-Pes.	2.26	[4.11]

Table 4.7
Momentum Profits Conditional on Investor Sentiment and Firm Size

This table presents average monthly returns in percentages for price momentum strategies involving all NYSE/AMEX stocks for the time period April 1967 until December 2008. Panel A shows momentum strategies implemented on the 50% smallest companies in the sample and Panel B in the 50% largest. Size is measured as price x shares outstanding at the end of the formation period. Size decile breakpoints are from Kenneth French's data library. We allow one month between the end of the formation period and the holding period, and delete all stocks that are priced less than one \$1 at the beginning of the holding period. Sentiment is defined in Table 2. In this table we use 30%-30% cut-off points for optimistic and pessimistic sentiment. To test whether momentum profits in each sentiment state respectively are equal to zero, we regress the time series of average monthly momentum profits on Optimistic, Pessimistic and Mild sentiment dummies, with no intercept. To test if mean profits in Optimistic sentiment periods are different from profits in Pessimistic sentiment periods we regress average monthly momentum profits on a Mild sentiment dummy variable and an Optimistic sentiment dummy variable with a constant. The *t*-statistics of the significance of momentum profits and the difference between profits derived after optimistic and pessimistic periods are calculated using Newey-West standard errors, where the lag is set to *K-1*. In this table *J*, *K*=6.

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
Panel A: Small Cap.													
Optimistic	(n=121)	-0.67	0.07	0.36	0.62	0.84	0.90	1.08	1.11	1.28	1.47	2.14	[6.16]
Mild	(n=286)	-0.48	0.20	0.54	0.74	0.89	0.99	1.06	1.16	1.24	1.23	1.72	[6.48]
Pessimistic	(n=93)	2.13	2.51	2.65	2.65	2.55	2.51	2.48	2.52	2.52	2.58	0.46	[0.88]
											Opt.-Pes.	1.68	[2.69]
Panel B: Large Cap.													
Optimistic	(n=121)	0.35	0.60	0.73	0.85	0.78	0.87	0.96	0.86	1.05	1.21	0.86	[2.17]
Mild	(n=286)	0.22	0.55	0.69	0.75	0.80	0.77	0.81	0.86	0.96	1.03	0.81	[3.05]
Pessimistic	(n=93)	1.80	1.87	1.87	1.91	1.80	1.64	1.67	1.62	1.72	2.07	0.26	[0.60]
											Opt.-Pes.	0.60	[1.01]

Table 4.8
Risk-adjusted Momentum Profits Conditional on Investor Sentiment

This table presents risk adjusted momentum profits calculated from CAPM, Fama-French and Conditional CAPM models. For each momentum portfolio and holding period month we form a time series of returns, which we regress on excess market return when we risk adjust according to the CAPM, and excess market return, the SMB and HML factors when we risk adjust according to the Fama-French 3 factor model. For the CCAPM we allow beta to differ depending on the average sentiment in the 6 formation periods that correspond to each portfolio return observation (see equation 3). Using these loadings and the factor realizations in each month, we estimate the monthly excess return for each portfolio. The data on market returns, the risk free rate and the SMB and HML factors are from Kenneth French's data library. Sentiment is defined as in Table 2. In this table we use 30%-30% cut-off points for optimistic and pessimistic sentiment. To test whether momentum profits in each sentiment state respectively are equal to zero, we regress the time series of average monthly momentum profits on Optimistic, Pessimistic and Mild sentiment dummies, with no intercept. To test if mean profits in Optimistic sentiment periods are different from profits in Pessimistic sentiment periods we regress average monthly momentum profits on a Mild sentiment dummy variable and an Optimistic sentiment dummy variable with a constant. The t-statistics of the significance of momentum profits and the difference between profits derived after optimistic and pessimistic periods are calculated using Newey-West standard errors, where the lag is set to $K-1$. In this table $J, K=6$.

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
Panel A: CAPM													
Optimistic	(n=121)	-0.83	-0.13	0.24	0.41	0.61	0.69	0.85	0.84	0.98	1.20	2.03	[5.46]
Mild	(n=286)	-0.52	0.14	0.44	0.60	0.69	0.75	0.82	0.87	0.93	0.97	1.49	[5.93]
Pessimistic	(n=93)	0.72	1.22	1.34	1.25	1.22	1.20	1.10	1.09	1.11	1.21	0.48	[1.10]
											Opt.-Pes.	1.55	[2.68]
Panel B: FF													
Optimistic	(n=121)	-0.87	-0.25	0.10	0.26	0.47	0.55	0.71	0.72	0.90	1.21	2.08	[5.39]
Mild	(n=286)	-0.82	-0.14	0.19	0.35	0.47	0.53	0.61	0.67	0.74	0.80	1.61	[6.83]
Pessimistic	(n=93)	-0.46	0.18	0.39	0.36	0.39	0.42	0.33	0.35	0.40	0.51	0.96	[2.30]
											Opt.-Pes.	1.12	[1.95]
Panel C: Conditional CAPM													
Optimistic	(n=121)	-0.83	-0.12	0.25	0.41	0.61	0.70	0.85	0.84	0.99	1.20	2.03	[5.46]
Mild	(n=286)	-0.52	0.13	0.43	0.59	0.68	0.74	0.81	0.87	0.92	0.96	1.48	[5.92]
Pessimistic	(n=93)	0.64	1.09	1.19	1.09	1.06	1.04	0.95	0.95	0.98	1.11	0.47	[1.07]
											Opt.-Pes.	1.56	[2.70]

Table 4.9
Momentum Profits Conditional on Sentiment Orthogonal to Current, Future Macroeconomic Conditions, and VIX

This table presents average monthly returns in percentages for price momentum strategies involving all NYSE/AMEX stocks for the time period January 1985 until December 2008. The description of the momentum strategy is defined in Table 2. Sentiment is measured using the time series of consumer confidence sentiment index constructed by Conference Board. We regress this series on growth in industrial production, real growth in durable, non-durable, and services consumption, growth in employment, an NBER recession indicator, one quarter ahead growth in industrial production, durable, non-durable, and services consumption, and employment and the closing level of the VIX at the last day of the month in which sentiment is measured. We use the residuals from this regression as the sentiment proxy. In this table we use 30%-30% cut-off points for optimistic and pessimistic sentiment. To test whether momentum profits in each sentiment state respectively are equal to zero, we regress the time series of average monthly momentum profits on Optimistic, Pessimistic and Mild sentiment dummies, with no intercept. To test if mean profits in Optimistic sentiment periods are different from profits in Pessimistic sentiment periods we regress average monthly momentum profits on a Mild sentiment dummy variable and an Optimistic sentiment dummy variable with a constant. The *t*-statistics of the significance of momentum profits and the difference between profits derived after optimistic and pessimistic periods are calculated using Newey-West standard errors, where the lag is set to *K-1*.

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
Panel A: J=6, K=3													
Optimistic	(n=67)	-0.95	0.01	0.56	0.77	0.81	0.90	1.11	0.94	1.11	1.33	2.27	[3.09]
Mild	(n=142)	-0.66	0.35	0.67	0.72	0.81	0.81	0.75	0.73	0.77	0.97	1.62	[3.59]
Pessimistic	(n=63)	1.04	1.20	1.28	1.29	1.36	1.26	1.21	1.08	1.24	1.40	0.36	[0.66]
											Opt.-Pes	1.91	[2.09]
Panel B: J, K=6													
Optimistic	(n=54)	-1.44	-0.47	0.06	0.27	0.48	0.52	0.72	0.68	0.84	0.98	2.42	[4.34]
Mild	(n=167)	-0.53	0.26	0.58	0.68	0.79	0.86	0.88	0.88	0.90	0.99	1.51	[4.73]
Pessimistic	(n=51)	1.88	1.88	1.82	1.73	1.69	1.66	1.61	1.67	1.82	2.15	0.27	[0.43]
											Opt.-Pes.	2.15	[2.52]
Panel C: J=6, K=12													
Optimistic	(n=74)	-0.43	0.12	0.39	0.46	0.58	0.64	0.67	0.70	0.70	0.70	1.13	[2.53]
Mild	(n=125)	-0.05	0.38	0.70	0.80	0.94	1.01	1.08	1.09	1.06	0.92	0.97	[3.22]
Pessimistic	(n=74)	1.05	1.25	1.23	1.14	1.15	1.10	1.06	1.10	1.19	1.26	0.21	[0.57]
											Opt.-Pes.	0.92	[1.61]

Table 4.10
Momentum Profits Conditional on an Alternative Investor Sentiment Index

This table presents average monthly returns in percentages for price momentum strategies involving all NYSE/AMEX stocks for the time period October 1965 until December 2007. The momentum strategy is defined in table 2. Sentiment is measured using the monthly sentiment index constructed by Baker and Wurgler (2007), using trading volume (measured as total NYSE turnover), dividend premium, closed-end fund discount, number and first day returns in IPO's, and the equity share in new issues. Because these variables are partly related to economic fundamentals, Baker and Wurgler regress each proxy against growth in industrial production, real growth in durable, non-durable, and services consumption, growth in employment, and an NBER recession indicator, and use the residuals from this regression as the sentiment proxies. The overall sentiment index is the first principal component of the six sentiment proxies. In order to identify whether a particular formation period was optimistic or pessimistic we follow the same procedure as that outlined in Table 2. In this table we use 30%-30% cut-off points for optimistic and pessimistic sentiment and we group the Mild sentiment and Optimistic sentiment categories together. To test whether momentum profits in each sentiment state respectively are equal to zero, we regress the time series of average monthly momentum profits on an Optimistic sentiment dummy variable and a Pessimistic sentiment dummy variable, with no intercept. To test if mean profits in Optimistic sentiment periods are different from profits in Pessimistic sentiment periods we regress average monthly momentum profits on an Optimistic sentiment dummy variable with a constant. The t-statistics of the significance of momentum profits and the difference between profits derived after optimistic and pessimistic periods are calculated using Newey-West standard errors, where the lag is set to $K-1$. In Panel A $K=3$, in Panel B $K=6$ and in Panel C $K=12$.

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
<i>Panel A: J=6, K=3</i>													
Optimistic	(n=371)	-0.24	0.54	0.85	0.96	1.04	1.09	1.14	1.12	1.27	1.46	1.71	[6.87]
Pessimistic	(n=136)	1.96	2.06	2.07	2.02	1.87	1.77	1.84	1.83	1.87	2.18	0.23	[0.46]
											Opt.-Pes.	1.48	[2.69]
<i>Panel B: J=6, K=6</i>													
Optimistic	(n=387)	-0.18	0.47	0.76	0.88	0.98	1.04	1.1	1.14	1.25	1.4	1.59	-0.18
Pessimistic	(n=120)	2.31	2.34	2.29	2.21	2.11	2.04	2.1	2.15	2.28	2.61	0.3	[0.21]
											Opt.-Pes.	1.29	[2.86]
<i>Panel C: J=6, K=12</i>													
Optimistic	(n=411)	0.21	0.62	0.84	0.91	1.01	1.05	1.09	1.12	1.15	1.17	0.97	[5.16]
Pessimistic	(n=96)	2.80	2.66	2.55	2.48	2.39	2.30	2.41	2.48	2.62	2.78	-0.01	[-0.02]
											Opt.-Pes.	0.98	[1.88]

Table 4.11
Momentum Profits Conditional on Different Specifications of Investor Sentiment

This table presents average monthly returns in percentages for price momentum strategies involving all NYSE/AMEX stocks for the time period April 1967 until December 2008. We allow one month between the end of the formation period and the holding period, and delete all stocks that are priced less than one \$1 at the beginning of the holding period. Sentiment is defined in Table 2. In this table we use 30%-30% cut-off points for optimistic and pessimistic sentiment. In order to identify whether a particular formation period was optimistic or pessimistic, in each month t we calculate the average sentiment level for the previous 2 (Panel A) and 4 (Panel B) months. To test whether momentum profits in each sentiment state respectively are equal to zero, we regress the time series of average monthly momentum profits on Optimistic, Pessimistic and Mild sentiment dummies, with no intercept. To test if mean profits in Optimistic sentiment periods are different from profits in Pessimistic sentiment periods we regress average monthly momentum profits on a Mild sentiment dummy variable and an Optimistic sentiment dummy variable with a constant. The t-statistics of the significance of momentum profits and the difference between profits derived after optimistic and pessimistic periods are calculated using Newey-West standard errors, where the lag is set to $K-1$. In this table $J, K=6$.

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.].
		<i>Panel A: Lag 2 sentiment</i>											
Optimistic	(n=118)	-0.46	0.19	0.50	0.66	0.84	0.93	1.10	1.10	1.27	1.49	1.95	[5.48]
Mild	(n=294)	-0.25	0.36	0.66	0.80	0.89	0.94	1.00	1.06	1.11	1.18	1.42	[5.73]
Pessimistic	(n=89)	2.14	2.45	2.47	2.32	2.26	2.24	2.16	2.19	2.32	2.59	0.46	[0.99]
											Opt.-Pes.	1.49	[2.56]
		<i>Panel B: Lag 4 sentiment</i>											
Optimistic	(n=122)	-0.43	0.23	0.60	0.76	0.95	1.03	1.20	1.18	1.37	1.66	2.10	[5.66]
Mild	(n=280)	-0.30	0.30	0.61	0.76	0.85	0.90	0.96	1.02	1.05	1.07	1.37	[5.25]
Pessimistic	(n=97)	1.97	2.32	2.33	2.19	2.13	2.12	2.04	2.08	2.23	2.59	0.62	[1.69]
											Opt.-Pes.	1.48	[2.83]

Table 4.12
Long-run Profits of Momentum Portfolios Conditional on Investor Sentiment

This table presents long-run event time returns for momentum portfolios formed after optimistic and pessimistic periods. J and K in this table are equal to 6. For each momentum portfolio we define an event period 13 months after the initial formation period. From this event date month onwards we estimate the average monthly return of this portfolio in the following 6 years. The final return of each portfolio is the geometric average of these monthly average profits. Panel A uses raw returns, Panel B CAPM adjusted returns and Panel C returns adjusted according to the Fama-French 3 factor model. Sentiment is defined as in Table 2. In this table we use 30%-30% cut-off points for optimistic and pessimistic sentiment. To test whether momentum profits in each sentiment state respectively are equal to zero, we regress the time series of average monthly momentum profits on Optimistic, Pessimistic and Mild sentiment dummies, with no intercept. To test if mean profits in Optimistic sentiment periods are different from profits in Pessimistic sentiment periods we regress average monthly momentum profits on a Mild sentiment dummy variable and an Optimistic sentiment dummy variable with a constant. The t-statistics of the significance of momentum profits and the difference between profits derived after optimistic and pessimistic periods are calculated using Newey-West standard errors, where the lag is set to the number of overlapping strategies, which is 6.

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
Panel A: 6 years holding period													
Panel A1: Raw													
Optimistic	(n=121)	1.04	0.93	0.87	0.84	0.81	0.77	0.74	0.70	0.63	0.48	-0.56	[-7.78]
Mild	(n=218)	1.39	1.40	1.39	1.38	1.37	1.36	1.35	1.33	1.31	1.22	-0.17	[-2.62]
Pessimistic	(n=82)	1.07	1.22	1.22	1.24	1.24	1.23	1.23	1.19	1.17	1.05	-0.02	[-0.18]
											Opt.-Pes.	-0.54	[-4.03]
Panel A2: CAPM													
Optimistic	(n=121)	1.05	0.91	0.84	0.81	0.77	0.74	0.70	0.67	0.61	0.48	-0.56	[-8.26]
Mild	(n=218)	0.89	0.94	0.95	0.94	0.94	0.93	0.91	0.89	0.84	0.72	-0.18	[-2.63]
Pessimistic	(n=82)	0.34	0.54	0.58	0.61	0.62	0.61	0.61	0.57	0.52	0.34	0.00	[0.03]
											Opt.-Pes.	-0.56	[-4.29]
Panel A3: FF													
Optimistic	(n=121)	0.63	0.49	0.43	0.41	0.39	0.37	0.35	0.33	0.29	0.18	-0.45	[-5.64]
Mild	(n=218)	0.42	0.49	0.51	0.53	0.55	0.55	0.56	0.54	0.51	0.40	-0.02	[-0.25]
Pessimistic	(n=82)	0.18	0.34	0.36	0.39	0.42	0.41	0.43	0.40	0.38	0.29	0.11	[1.16]
											Opt.-Pes.	-0.56	[-4.50]

Table 4.12. Continued

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
		Panel B: 5 years holding period											
Panel B1: Raw													
Optimistic	(n=121)	0.92	0.82	0.77	0.74	0.71	0.67	0.64	0.60	0.52	0.37	-0.55	[-6.82]
Mild	(n=228)	1.39	1.38	1.38	1.36	1.36	1.34	1.34	1.33	1.30	1.21	-0.18	[-2.30]
Pessimistic	(n=84)	1.13	1.28	1.29	1.31	1.32	1.30	1.30	1.27	1.22	1.08	-0.05	[-0.37]
											Opt.-Pes.	-0.50	[-3.36]
Panel B:2 CAPM													
Optimistic	(n=121)	1.02	0.89	0.82	0.78	0.75	0.72	0.69	0.65	0.59	0.47	-0.55	[-7.15]
Mild	(n=228)	0.84	0.88	0.90	0.90	0.90	0.88	0.87	0.85	0.80	0.66	-0.18	[-2.28]
Pessimistic	(n=84)	0.34	0.56	0.60	0.64	0.66	0.65	0.64	0.60	0.53	0.32	-0.02	[-0.16]
											Opt.-Pes.	-0.53	[-3.56]
Panel B3: FF													
Optimistic	(n=121)	0.63	0.49	0.42	0.40	0.39	0.36	0.35	0.33	0.29	0.20	-0.43	[-4.78]
Mild	(n=228)	0.40	0.45	0.49	0.51	0.54	0.53	0.54	0.53	0.50	0.40	0.00	[-0.01]
Pessimistic	(n=84)	0.10	0.30	0.33	0.38	0.41	0.40	0.42	0.40	0.36	0.22	0.12	[1.07]
											Opt.-Pes.	-0.55	[-3.82]

Table 4.12. Continued.

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
		Panel C: 4 years holding period											
Panel C1: Raw													
Optimistic	(n=121)	0.81	0.73	0.70	0.68	0.67	0.65	0.62	0.57	0.49	0.32	-0.49	[-5.79]
Mild	(n=233)	1.36	1.36	1.37	1.35	1.36	1.34	1.34	1.32	1.30	1.20	-0.17	[-1.86]
Pessimistic	(n=91)	1.13	1.28	1.29	1.29	1.30	1.27	1.27	1.24	1.20	1.06	-0.06	[-0.54]
											Opt.-Pes.	-0.43	[-2.90]
Panel C2: CAPM													
Optimistic	(n=121)	0.96	0.85	0.80	0.77	0.76	0.73	0.71	0.66	0.60	0.48	-0.49	[-6.26]
Mild	(n=233)	0.78	0.83	0.86	0.86	0.87	0.85	0.84	0.82	0.77	0.62	-0.17	[-1.83]
Pessimistic	(n=91)	0.37	0.59	0.63	0.65	0.67	0.65	0.65	0.60	0.53	0.33	-0.04	[-0.33]
											Opt.-Pes.	-0.45	[-3.12]
Panel C3: FF													
Optimistic	(n=121)	0.60	0.47	0.42	0.41	0.41	0.39	0.38	0.35	0.31	0.21	-0.38	[-4.34]
Mild	(n=233)	0.36	0.42	0.48	0.49	0.52	0.51	0.52	0.51	0.48	0.37	0.01	[0.14]
Pessimistic	(n=91)	0.10	0.29	0.33	0.36	0.39	0.38	0.40	0.37	0.33	0.22	0.13	[1.13]
											Opt.-Pes.	-0.51	[-3.58]

Table 4.12. Continued

		Momentum Portfolio											
		1=Sell	2	3	4	5	6	7	8	9	10=Buy	Buy-Sell	[t-stat.]
		Panel D: 3 years holding period											
Panel D1: Raw													
Optimistic	(n=121)	0.69	0.66	0.66	0.66	0.67	0.65	0.62	0.58	0.49	0.32	-0.37	[-3.62]
Mild	(n=245)	1.31	1.29	1.31	1.29	1.30	1.29	1.29	1.29	1.27	1.15	-0.16	[-1.66]
Pessimistic	(n=91)	1.05	1.25	1.28	1.28	1.29	1.26	1.27	1.25	1.20	1.07	0.02	[0.08]
											Opt.-Pes.	-0.39	[-1.81]
Panel D2: CAPM													
Optimistic	(n=121)	0.84	0.77	0.75	0.74	0.75	0.72	0.70	0.67	0.60	0.47	-0.38	[-3.88]
Mild	(n=245)	0.73	0.75	0.80	0.80	0.81	0.80	0.80	0.78	0.73	0.57	-0.16	[-1.65]
Pessimistic	(n=91)	0.37	0.62	0.69	0.70	0.72	0.70	0.71	0.68	0.60	0.41	0.05	[0.24]
											Opt.-Pes.	-0.43	[-1.97]
Panel D3: FF													
Optimistic	(n=121)	0.57	0.46	0.43	0.43	0.46	0.44	0.43	0.42	0.36	0.28	-0.29	[-2.82]
Mild	(n=245)	0.31	0.35	0.42	0.44	0.48	0.49	0.50	0.50	0.47	0.35	0.04	[0.46]
Pessimistic	(n=91)	-0.03	0.22	0.30	0.33	0.38	0.38	0.40	0.39	0.35	0.24	0.27	[1.44]
											Opt.-Pes.	-0.56	[-2.61]

5. Decision heuristics and Bayesian updating

5.1 Introduction

Bayesian updating is at the core of rational decision making. Various psychological studies present evidence that people systematically violate this paradigm. The economic implications of this evidence can be severe. For example, a maintained assumption in neoclassical models of asset prices, such as the CAPM or the APT, is that investors update their expectations of assets' risk returns profiles in a Bayesian manner. However, the evidence presented by psychologists rejects this assumption and points to the existence of systematic mispricings in the marketplace, as expectations, and hence prices, are updated in a biased manner. Research in behavioural finance implements these findings from psychology in models of asset prices in an attempt to explain the apparent anomalous behaviour of asset prices [see for example Daniel et al (1998), Odean (1998) and Barberis, Shleifer and Vishny (1998)].

However, the validity of these models crucially depends on the robustness of the evidence of non-Bayesian updating. One view held by experimental economists, is that this evidence of non Bayesian updating is debatable, because they are generally documented in experiments which *do not* comply with certain methodological requirements that have been documented to affect subjects' behaviour. However, because these anomalies are potentially relevant to economic theory, economists seek to examine their validity in experiments that satisfy these requirements [see for example Grether and Plott (1979)]. In this spirit we report results from economics experiments designed to test the validity of a behavioural theory of non-Bayesian updating proposed by Griffin and Tversky (1992), which has been used to model the behaviour of asset prices [Barberis, Shleifer and Vishny (1998)].

Griffin and Tversky (1992) (hereafter GT) find that people over-predict probabilities when the information available is of high strength (salience) and low weight (predictive validity/credence), and under-predict when strength is low and weight is high. They use the following example to explain what they mean by these words: Suppose you receive a recommendation letter for a recent PhD graduate written by an assistant professor stating that the graduate is outstanding. This letter is of high strength because its tone is warm. However, its weight is low because the assistant professor is inexperienced in assessing the future potential of graduates. Conversely, suppose that a highly experienced professor writes the letter saying that the student's research potential is satisfactory. In this case the strength of the letter is lower as the tone is less enthusiastic about the candidate. However, because the professor has seen many graduates in his career, the weight of the letter is high. GT hypothesize that people focus excessively on the strength of the information available and not as much as they should on the weight, resulting in overconfidence (underconfidence) when the former is high (low) and the latter low (high).

This property of the GT hypothesis is particularly appealing to finance because it can explain the (puzzling) positive and negative serial correlation in prices. That is, when the information is of low strength and high weight, investors will underreact inducing positive correlation. Conversely, when the information is of high strength and low weight investors will overreact, inducing negative correlation. This appealing property of the GT hypothesis has been utilized by asset pricing theories such as Barberis et al (1998) and Sorescu and Subrahmanyam (2006).

However, because the experiments reported by GT do not fulfil the methodological requirements of a controlled experiment advanced by economists, their robustness is debatable. This makes the economic applications less relevant. Specifically, three features of the GT

experiment can be seen as problematic: subjects were not incentivised to make correct decisions, the task involved responses to hypothetical and unnatural events, and subjects were implicitly assumed to be risk neutral throughout.

It is often argued that the design of incentives in experiments should encourage subjects to be accurate. That is, the payment of subjects should be directly related to their choices, and, whenever possible, the potential payoff from making the right choice should exceed the cost of making it [Smith (1982)]. These requirements are proposed because various experiments show that behaviour can change when accuracy is encouraged with monetary incentives [e.g., see reviews by Harrison (2006) and Harrison and Rutström (2008)]. A notable example is Grether (1980), who conducts an experiment to test whether agents form probabilities according to the representativeness heuristic, identified by Kahneman and Tversky (1972, 1973) in experiments that did not provide financial incentives. He concludes that the presence of incentives makes subjects behave as Bayesians. In the experiments designed by GT subjects' payments were independent of their choices. Given that "hypothetical bias" has been found to reduce behavioural anomalies, it is not clear whether their findings are robust to the presence of financial incentives.

Further, economic decisions are based on naturally-occurring, observable events. For example, investors in stock markets use information such as the recent earnings performance of companies and elect whether to purchase the stock. In insurance markets people use the available information to assess the likelihood of certain events, such as floods, in order to decide whether to buy insurance. This information can be the occurrence of past floods, the location of their residence and prevailing weather conditions. In these cases the information used is natural in the sense that people have experienced it occurring. However, in the design of GT subjects are presented with information about a *hypothetical* event and asked to use this information to

estimate a probability. The hypothetical event is that a biased coin is tossed a number of times, with a particular outcome emerging. Based on this information subjects are asked to estimate the probability of the bias favouring heads or tails. This design can be seen as unnatural because the event in question did not actually occur; therefore we cannot be certain that the process of “visualising” it does not, in and of itself, affect responses. In addition the task itself is difficult to grasp because subjects are asked to visualize a biased coin and evaluate a probability in terms of *the bias favouring either side*, not of the coin landing heads or tails. Arguably, this can confuse subjects since it is difficult to imagine such a coin. Also, the fact that subjects were explicitly asked to express a probability is itself problematic because the term has been debated amongst statisticians for years.

Finally, GT implicitly assume risk neutrality throughout. This is a very strong assumption, because virtually all experiments designed to estimate peoples’ risk preferences conclude that subjects are generally risk-averse [see Harrison and Rutström (2008) for a review]. In the presence of risk aversion, stated or revealed probabilities will be distorted by the impact of risk attitude on the expected payoff, therefore will not correspond to true beliefs [Savage (1971), Kaldane and Winkler (1988)]. Since GT implicitly assume risk neutrality throughout, we do not know whether their findings are robust to different risk preference specifications.

Despite these methodological concerns with the evidence for the hypothesis advanced by GT, researchers have used it to construct economic models [Barberis et al. (1998), Sorescu and Subrahmanyam (2006)]. This chapter contributes to the literature by taking a step back and examining whether this hypothesis holds when the previous methodological shortcomings are addressed. In this manner we test whether it is robust as an alternative decision-making model to Bayesian updating.

We implement an incentive compatible design, where subjects are encouraged to provide accurate responses. We rely solely on natural frequencies as the information used by subjects is generated in the experiment. We avoid explicitly asking subjects to estimate a probability, and instead implement a betting task that allows extraction of subjective beliefs, as in Fiore et al (2008). Finally, we go beyond risk neutrality and estimate probabilities allowing for risk aversion.

Our results broadly support the original findings of GT. We find that when the information presented to subjects is of high strength and low weight, elicited probabilities are higher, compared to instances that the information is of low strength and high weight. Specifically, we find that strength outweighs the effect of weight on judgment roughly by a factor of 2. Since, according to Bayes Rule, strength and weight should affect judgment equally, our findings suggest that the strength-weight hypothesis captures the process with which beliefs are updated more accurately than Bayesian Updating.

However, one important difference between our results and those reported by GT is that we find a much smaller spread between elicited probabilities from information of high strength and low weight, and elicited probabilities from information of low strength and high weight. For example, GT in one particular instance find an alarmingly large spread of 25% between probabilities that *should* be equal. In our analysis the corresponding figure is 4.7%. Of course comparing point estimates between different experiments is not ideal, but we suggest that the implementation of an incentive compatible design, the use of naturally occurring information, and controlling for risk attitudes, reduce the behavioural effects identified by GT.

Apart from the methodological contributions already highlighted the systematic deviations from Bayesian Updating found in this chapter challenge the argument used by Fama (1998) that market prices are formed *as if* investors update their expectations using Bayes law [in

the spirit of Friedman (1952)]. Rather, our findings show that decision makers *as a group* behave in a manner that leads to biased expectations. Since arbitrage forces are constrained (see discussion in section 2.4), the marketplace implication of this finding is that prices will be set according to the strength and weight of the information set, and not according to fundamental values. Therefore, our analysis, by establishing the robustness of the psychological mechanism that triggers biased expectations initially proposed by Griffin and Tversky (1992), increases our understanding about the price formation process.

Section 2 explains the experimental design, and the way we operationalize strength and weight. Section 3 explains how we jointly estimate risk attitudes and subjective probabilities, and presents the results from this joint estimation. It also explains in detail the model we use to estimate the exact effect of strength and weight on elicited beliefs. Section 4 concludes the chapter.

5. 2 Experimental Design

The format of our experiment was chosen so that we embed the elicitation procedure in tasks that allow us to elicit subjects' beliefs about events that differ in strength and weight, train subjects to understand the belief elicitation technique, and identify attitudes towards risk. In no particular order, we performed two tasks in this experiment. The first task was aimed at identifying risk attitudes, and the second to elicit their subjective beliefs for events that differ in strength and weight.

5.2.2 Characterizing Attitudes Towards Risk

The first series of tasks were designed to reveal attitude towards risk. This information will be useful when we estimate subjects' subjective beliefs allowing for risk aversion. For this purpose we use a series of binary choice tasks, following Hey and Orme (1994). Each lottery consists of one, two or three monetary prizes, with four possible monetary values, £0, £5, £10 or £15. Table 5.1 displays a typical lottery pair. In our experiment we divided these 60 tasks into 3 groups of 20, and each subject made choices for one of these groups. At the end of the task one of the twenty choices of each subject was selected at random and played out for money. Appendix 1 shows the instructions given to subjects for this particular task.

5.2.3 Eliciting Subjective Beliefs

The second task of our experiment is aimed at eliciting beliefs for events of different strength and weight. In designing these events we allow them to occur randomly. The task involves two boxes, one white and one blue, each containing ten-sided dice, shown in Figure 5.1. The white box contains dice with six white sides and four blue sides, and the blue box contains dice with six blue sides and four white sides.⁶² The number of dice in the boxes varies in the experiment, but the two boxes always contain the same number of dice. In the beginning of each round we roll a six-sided dice, with three blue sides and three white sides. If a blue (white) side comes up, we roll all the dice in the blue (white) box and announce the outcome. For example, suppose that the two boxes contain 3 dice each, and that after we roll the six-sided die we get a white face. This means that we roll all the dice in the white box and announce the outcome; for example, 2 dice showing a white face and 1 dice showing blue. This information is announced

⁶² We use 0.6 and 0.4 because these are the probabilities used by GT. Figure 5.1 shows pictures of the boxes and dice.

and written on a whiteboard. Because the process of rolling the six-sided die and the dice in the chosen box takes place behind a screen, the task for subjects is to consider the information provided and update their belief of the white box being chosen initially.

With no additional information from the rolling of dice from the chosen box, the probability of either box being used is 0.5. By announcing the outcome, this prior should be updated in accordance with Bayes Rule. GT suggest that, contrary to Bayes Rule, individuals are more sensitive to the strength of the evidence (how many white or blue sides came up as a proportion of the total number of dice that were rolled) as opposed to the weight (how many dice were rolled in total). In our experiment we test this conjecture by comparing subjects' betting patterns *for dice patterns that have the same posterior but differ in strength and weight*. For example, if we roll 5 dice and get 4 dice with a white face and 1 dice with a blue face, the posterior probability of the white box being used is the same (0.77) as when we toss 9 dice and get 6 white and 3 blue. Bayes Rule predicts that the elicited probabilities in these two cases should be identical. However, GT predict that, because the strength of the former pattern is greater, $3/5 > 3/9$, and the weight smaller, $5 < 9$, beliefs will be updated more strongly when we get 4 white and 1 blue as opposed to 6 white and 3 blue.

When the outcome from rolling all the dice in the chosen box is announced, subjects are asked to consider the chances that this box is the blue or the white box, and place a series of bets. From the betting behaviour of each subject in each round we can deduce their subjective belief of either box being used. They are given a £3 stake with which they place a bet in each of 19 different betting houses that give different odds on the white or blue box being chosen. The £3 from one bookie is not transferable to another bookie. Since the probabilities that the bookies assign to the two events are the inverse of the odds, assuming no house take, these bookies are essentially placing different probabilities on the white or the blue box being used.

Table 5.2 shows the betting sheet where subjects are asked to place their bets. For each bookie (row), subjects consider the odds offered and their subjective probability for the two events, i.e., the blue or the white box being chosen, and place their £3 stake on either the blue or the white box for each bookie. The illustration in Table 5.2 shows a subject that was willing to bet on the white box until the fourth bookie, which offers 5 to 1 odds that the white box is being used, and switches to betting on the blue box from bookie 5 onwards. Assuming no house take the probability assigned to the event that the dice come from the white box from the 4th bookie is the inverse of the odds, i.e., $1/5=0.20$, and from the 5th bookie $1/4=0.25$. Thus, assuming risk neutrality, we have trapped the subjective probability of the subject in the closed interval $[0.20, 0.25]$.

The intuition of this design is that, when the probability for the white box offered is lower than 0.20, the odds offered make the expected value of a bet on the white box, i.e., the subjective probability multiplied by the payout, greater than that of a bet on the blue box. Therefore, the subject bets on the white box. However, when the offered probability exceeds the subject's subjective belief, the expected value of a bet on the white box is less than that of a bet on the blue, therefore the subject switches bet. A detailed example of the theoretical foundation of this method is given in the next section. This method of belief elicitation has been developed by Andersen, Fountain, Harrison and Rutström (2007), and assumes that subjects are capable of forming subjective beliefs in the sense of Machina and Schmeidler (1992; p. 747). Such decision makers behave as if they employ subjective probabilities of events and utilities of outcomes that are independent of the assignment of outcomes to events.⁶³

⁶³ This requirement is not strong in our data because a violation of probabilistic sophistication arises in conditions of ambiguity, which depend on whether relevant information about the probabilities of interest is missing (see Ellsberg 1961). In our design *all* the relevant information is revealed, therefore ambiguity does not arise.

The betting process is repeated 30 times in each session for each subject. We play 4 rounds where the boxes contain 3 dice, 14 rounds where the boxes contain 5 dice, 6 rounds with 9 dice, and 6 rounds with 17 dice. This distribution of sample sizes was chosen to ensure that we observe roughly the same amount of “extreme” samples.⁶⁴ A subject from each session was chosen randomly to act as a monitor. The monitor was in charge of rolling and counting the dice, as well as announcing the outcome.

We performed 12 sessions in total. In order to control for possible framing effects, in half of the sessions we conducted the risk task first, followed by the belief task. In the belief task, we presented the sample sizes in ascending order (i.e., first 4 rounds of 3, then 14 rounds of 5, etc) and in descending order (6 rounds of 17, 6 rounds of 9, etc). Therefore we have an overall 2x2 design, with 4 treatments in total.

Our instructions illustrate the factors that affect betting in real prediction markets, such as betting on a horse race with different bookies, and drawing parallels between such naturally occurring events and our task. Subjects first read these instructions quietly; we then read them aloud and allow time for questions. These instructions are shown in Appendix 2. We play 3 practice rounds, with the boxes containing 4 dice, prior to starting the real betting task, so that subjects become accustomed to the process. In order to determine payment to the subject at the end of the 30 rounds we randomly choose one particular round and bookie and play that bet for real. For example, each subject filled out 30 betting sheets, such as the one shown by Table 5.2. At the end we first randomly choose for each subject one of these betting sheets, and then 1 of the 19 bookies within that selected sheet. If, for that particular bookie, the subject placed the

⁶⁴ For example, the patterns that are the most difficult to get are (5,0) with probability of occurring equal to 0.044, pattern (7,2) with probability 0.091, and pattern (11,6) with probability 0.104. Based on these probabilities we chose the frequency of each sample size so that we roughly equalize their expected occurrence in each session.

allocated £3 on the box that was actually chosen, he got paid the amount that corresponds to the odds offered by that bookie.

We recruited 111 subjects from the University of Durham, U.K. Subjects were recruited using the computerized *ExLab* interface (<http://exlab.bus.ucf.edu>), after being solicited in general terms to register with *ExLab* for paid experiments. All subjects received a £5 show up fee. Apart from the tasks described above, each subject completed a survey of demographic characteristics, shown in Appendix 3. Payments for the experiment totalled £2,692 an average payment of £24.26 per subject.

5.2.4 Estimating the exact effect of strength and weight on elicited beliefs

In order to illustrate the model with which we estimate the effect of strength and weight on elicited beliefs, we use an example concerning two hypotheses (similar to the white or blue box being chosen) A and B where $P(A) + P(B) = 1$. In this case A can be that the white box is chosen, which includes dice with 0.6 probability of a white side and 0.4 probability of a blue side, and B can be that the blue box is chosen, which contains dice with 0.4 probability of a white side and 0.6 probability of a blue side. The task therefore is to determine the likelihood of A in comparison to the likelihood of B . Ceteris paribus our expectation of A is $P(A)$ and of B , $P(B)$. These probabilities in a Bayesian framework are called the priors.

Suppose that we observe an information signal C . In our case C is the pattern that emerged, (i.e., w dice showing a white face and $n-w$ dice showing a blue face) from tossing the dice in the chosen box at the beginning of each round. Bayes rule is the method with the priors should be combined with the new data, to get a posterior belief in terms of the hypotheses A and B . In other words the rule reveals how the occurrence of C should influence our expectation of A and B .

Formally the rule states that:

$$P(A/C) = \frac{P(C/A)P(A)}{P(C)} \quad (5.1)$$

Then if we divide the above with the corresponding probability of the likelihood of $P(B/C)$ we get:

$$\frac{P(A/C)}{P(B/C)} = \frac{P(C/A)}{P(C/B)} \quad (5.2)$$

Equation (5.2) shows the likelihood ratio. Each hypothesis, therefore, is distributed binomially with parameters n, p . In terms of our experiment, C is the sample of n dice that we rolled from the chosen box. The two competing hypotheses are that the dice came from the white box or the blue box.

$$P(C/A) = \binom{n}{w} 0.6^w 0.4^{n-w} \quad (5.3)$$

Equation (5.3) yields the probability of W white dice and $n-w$ blue dice, from the white box. $P(B/C)$ is the corresponding probability of the particular pattern coming from the blue box.

Dividing $P(C/A)$ by $P(C/B)$ gives:

$$\frac{P(A/C)}{P(B/C)} = \frac{P(C/A)}{P(C/B)} = \frac{\binom{n}{w} 0.6^w 0.4^{n-w}}{\binom{n}{w} 0.6^{n-w} 0.4^w} = \frac{0.6^{-n+2w}}{0.4^{-n+2w}} \quad (5.4)$$

Taking logs of both sides:

$$\text{Log}\left(\frac{P(A/C)}{P(B/C)}\right) = (-n + 2w)\log\left(\frac{0.6}{0.4}\right) \quad (5.5)$$

Because $n = W + B$, substituting for $-n$ yields:

$$\text{Log}\left(\frac{P(A/C)}{P(B/C)}\right) = (w - b)\log\left(\frac{0.6}{0.4}\right) \quad (5.6)$$

Now multiply and divide by n yields:

$$\text{Log}\left(\frac{P(A/C)}{P(B/C)}\right) = n \frac{(w - b)}{n} \log\left(\frac{0.6}{0.4}\right) \quad (5.7)$$

The right hand side is the probabilities elicited from subjects' responses in the belief task modified to log-odds form. If they are Bayesian they should be explained by the right hand side, which was derived from Baye's Theorem. The two dimensions that GT suggest are sample size (weight), n , and strength $(w - b)/n$, i.e., the amount of dice showing a white face minus the amount of dice showing a blue face, all divided by the total amount of dice rolled. In order to decompose we take logs of both sides, which yields:

$$\text{Log}\left[\frac{\text{Log}\left(\frac{P(A/C)}{P(B/C)}\right)}{\text{Log}\left(\frac{0.6}{0.4}\right)}\right] = a \log(n) + \beta \log\left(\frac{(w - b)}{n}\right) \quad (5.8)$$

The coefficients a and β that are derived from the above regression are in effect the importance that the individuals have given to strength and weight. Under the null the coefficients should be

equal to 1, because Bayes Rule predicts that they should affect judgment equally. Under the GT hypothesis $\alpha > \beta$.

5.3 Results

5.3.1 Elicited Beliefs From The Raw Data

We present the results of our experiment in two stages. First we analyze the raw responses from our belief task, which can be interpreted as elicited beliefs under the assumption of risk neutrality. Of course these beliefs will be biased if subjects are in fact risk averse or risk seeking. However, we explicitly show the elicited beliefs under the assumption of risk neutrality in order to highlight the effect of risk attitude on elicited beliefs, and how inferences crucially depend on the assumptions one makes about the shape of the utility function. In a second stage we relax this assumption, and define a joint likelihood over the responses to the risk and the belief task, allowing us to generate maximum likelihood estimates of both risk aversion coefficients and subjective beliefs.

Table 5.3 shows the composition of the different dice patterns, their strength and weight characteristics, the Bayesian probability of the white box given their occurrence, and the number of times subjects placed bets for each of them. We group together the patterns that have the same posterior probability, but differ in their strength and weight characteristics. Unless otherwise stated, when we refer to elicited or posterior probability we mean the probability of the white box being used. We have 6 such groups with a posterior probability of 0.12, 0.23, 0.4, 0.6, 0.77

and 0.88. The patterns that yield a posterior of 0.12 are (0, 5),⁶⁵ (2, 7) and (6, 11), the patterns that yield a posterior of 0.23 are (0, 3), (1, 4), (3, 6) and (7, 10), the patterns that yield a posterior of 0.4 are ((1, 2), (2, 3), (4, 5) and (8, 9), and *vice versa* for posteriors of 0.6, 0.77 and 0.88.⁶⁶ For example suppose that the boxes contain 5 dice, and after we roll the dice from the chosen box we observe 4 dice showing a white face and 1 dice showing a blue face (call this information “data”). The conditional probability that the box chosen initially was the white is:

$$P(\text{White} / \text{Pattern}) = \frac{P(\text{Data} / \text{White})P(\text{White})}{P(\text{Data})} \quad (5.9)$$

Substituting:

$$P(\text{White} / \text{Data}) = \frac{0.26 * 0.5}{0.17} = 0.77 \quad (5.10)$$

$P(\text{Data}/\text{White})$ is the probability that we obtain (with replacement) 4 dice showing a white face and 1 showing a blue face from the white box. $P(\text{White})$ is the probability that we use the white box, i.e, the prior, and $P(\text{data})$ is the probability that observe this particular data from either box.

If subjects update their beliefs using Bayes Rule, elicited beliefs for each pattern within each of the groups with the same posterior should be equal. However, under the hypothesis proposed by GT, we should observe that probabilities are higher when strength is high and weight is low.

⁶⁵ When we refer to a pattern of (3,0) we mean rolling 3 dice and getting 3 dice showing a white side and 0 dice showing a blue side. Similarly (1,4) means rolling 5 dice, where 1 dice shows a white face and 4 dice show a blue face.

⁶⁶ We observed 6 instances with a pattern of (13,4), 4 instances with a pattern of (12,5), 1 instance with a pattern of (14,3), 1 instance with a pattern of with (9,0), 1 instance with a pattern of (8,1), 4 instances with a pattern of (5,12), 4 instances with a pattern of (4,13), 1 instance with a pattern of (3,14), 2 instances with a pattern of (1,8) and 1 instance with a pattern of (2,15). These 25 pattern compositions do not appear in the GT study so we do not include them in our tables. However, we include them in the analysis to estimate the probabilities. Excluding them altogether from the analysis does not change any of our findings.

Figure 5.2 shows the distribution of raw elicited beliefs for each pattern when we group the patterns by posterior probability. We calculate raw elicited beliefs using subjects switch points. To illustrate, Table 5.2 shows the betting behaviour of a particular subject who switches to betting on the blue box from bookie 5 onwards. Assuming that the subject is risk neutral, this means that the subjective belief of this subject of the white box being used lies between 0.20 and 0.25, therefore the midpoint equals $(0.20 + 0.25)/2$.⁶⁷ We perform this calculation for all subjects in our sample for the particular dice pattern, and plot the distribution of raw elicited beliefs in Figure 5.2. The vertical bar in each panel shows the probability implied by applying Bayes rule without error.

The raw elicited beliefs in groups of patterns with equal posterior probability do not appear to be equal, at least under the maintained assumption that subjects are risk neutral. To illustrate, Figure 5.2 shows the distribution of switch points for patterns with 5, 9 and 17 dice rolled that have a posterior probability of 0.88, namely (5, 0), (7, 2) and (11, 6). It appears that subjects tend to get the probability right when the sample is (5, 0), and underestimate the probability as the sample size (i.e., the weight) increases to 9 and 17. This finding suggests that subjects, under the assumption of risk neutrality, update their beliefs more strongly when the pattern is of high strength and low weight, despite the fact that Bayes rule indicates that the probabilities in these cases are equal.

The top right Panel of Figure 5.2 shows the distribution of switch points for patterns with 5, 9 and 17 dice rolled that have a posterior probability of 0.12, namely (0, 5), (2, 7) and (6, 11). This figure is the mirror image of Figure 1.1. The elicited probability of the blue box, $1 -$ elicited probability of the white, is larger when the pattern is (0, 5) and decreases for patterns (2, 7) and (6, 11).

⁶⁷ If a subject bets on the white box for all the bookies his subjective belief lies in the interval $(0.95 + 1)/2$. If the subject bets on the blue box for all bookies his subjective belief is $(0.05)/2$.

Similar findings can be seen in the remainder of figure 5.2. Generally, subjects tend to be close to Bayesian probabilities when the pattern is of high strength and low weight, and underestimate those probabilities as the sample size (i.e., weight) decreases. An interesting finding is that the strength-weight effect on raw risk neutral beliefs appears to decline as the posterior probability moves towards 0.5. For example, the left-middle and left-bottom panels of Figure 5.2 show that the distributions of the switch points for patterns with posterior probabilities 0.6 and 0.4 do not significantly diverge. This suggests that the strength-weight effect is more influential when the patterns require a large update of the prior.

In order to formally derive an average subjective probability for each pattern we use interval regression techniques. The dependent variable is the interval which contains the belief of each subject in each round, and the independent variables are dummies for the various dice patterns, dummies for the order of the risk and the belief task ($task=1$ if risk task was conducted first) and the order with which the various sample sizes ($n=3, 5, 9, 17$) were presented in the belief task ($Descending=1$ in descending order).

Table 5.4 shows the results from the estimation. From the 22 patterns that we elicit probabilities, the hypothesis that they formed according to Bayes Rule is accepted at the 1% confidence interval for 9 cases. These cases mainly correspond to dice patterns that are of high strength and low weight, or when the Bayesian probability is close to 0.5. In order to compare the effects of strength and weight on the elicited probabilities we split them into 6 groups of equal posterior. In all 6 groups the hypothesis that the probabilities of the different patterns are equal is rejected at the 1% level of confidence. Rather, the results show that elicited probabilities are higher when strength is high and weight is low. To illustrate, consider the group with posterior probability 0.77. The elicited probability decreases monotonically from an average of

0.773 for the high strength low weight sample (3,0) to an average of 0.659 for the low strength high weight sample (10,7), a difference of 12.6 percentage points.

The first three panels in Table 5.4 correspond to dice patterns that point to the blue box being more likely, for example (0,3), (1,4) etc. To make direct comparisons with the elicited probabilities in the bottom 3 panels, one must subtract these probabilities from 1. For example, for the group of patterns with posterior probability of 0.12, the elicited probability for the high strength, low weight pattern (0,5) is $1 - 0.122 = 0.878$ and for the low strength, high weight sample (6,11) $1 - 0.250 = 0.75$.

The order and ascending/descending dummies show that elicited probabilities are higher by roughly 1.4 percentage points (p-value 0.043) when the belief task is conducted after the risk task, and by 1.9 percentage points (p-value 0.009) when the dice are presented in descending order. This suggests that the design of the experiment induces framing effects, although these are not substantial enough quantitatively.

5. 3.2 Joint Estimates of Beliefs and Preferences Towards Risk

In the previous section we elicit unconditional estimates of subjective probabilities assuming that our subjects are risk neutral. If they are risk averse or risk seeking the probabilities we elicit will not correspond to their true beliefs. Consistent with the general evidence from comparable experiments, we expect subjects will generally exhibit modest risk aversion. In this case elicited probabilities will be larger than true beliefs because, as shown by Andersen et al. (2007), a risk-averse individual will continue betting on her favoured alternative (i.e. the blue or the white box being used) beyond the bookie that corresponds to her subjective belief. This is because the subject's utility function is concave in the payoffs, so a larger expected payout from

the *least* favoured alternative is required to make the subject switch and bet on it. To illustrate, suppose we have two subjects, one risk averse and one risk neutral, who, after observing a particular dice pattern, believe that the probability of the white box being used is 0.75. Assume that the risk averse subject has a CRRA utility function defined over the gain domain of payouts y as:

$$U(y) = y^{1-r} / 1 - r \quad (5.11)$$

where r is the coefficient of risk aversion and $y \geq 0$. With this specification $r = 0$ implies risk neutrality, $r > 0$ implies risk aversion and $r < 0$ implies risk seeking. Table 5.5 shows how these two individuals place their bets on the white or the blue box for each bookie, given their belief, the odds offered by each bookie, and their preferences over risk. Columns 1 and 2 of Table 5.5 show the payout that each bookie offers for either the white or the blue box. The risk neutral person cares about the expected value (EV)⁶⁸ of the bet, and will bet on the white if $EV(W) > EV(B)$, as shown by columns 4 and 5. This risk neutral person will bet on the white box for bookies 1-15 and switch to betting on the blue box for all remaining bookies, as shown by the column BET RN. The risk averse person cares about expected utility (EU)⁶⁹ and will bet on the white box if $EU(W) > EU(B)$.⁷⁰ Because this person's utility function is concave in the payoffs, he continues betting on the white box past the bookie that corresponds to his belief, and switches to betting on the blue box only for the last bookie, as shown by column BET RA. In general, the

⁶⁸ For example the EV of a bet on the white for a bookie that pays £30 for a £3 pounds stake for an individual that believes that the probability of the white 0.6 is 0.6×30 .

⁶⁹ The EU of a bet on the white for a particular bookie is the probability weighted utility of the payout the bookie offers for a bet on the white. For example suppose that the bookie offers £30 for a £3 bet on the white, and that the subject believes that the probability of the white is 0.6. The EU of this bet is $0.6 \times (30)^{1-r} / 1-r$, where r is the subject's coefficient of risk aversion.

⁷⁰ The calculations of expected utility in Table 6 assume that the risk averse subject has risk aversion coefficient of 0.57. This is the average predicted value from experiments reported in Harrison, Johnson and, McInnes and Rutström (2005), and is consistent with estimates obtained by many other experimenters with comparable samples and stakes (e.g., Holt and Laury 2002).

more risk averse the subject is, the more delayed his switch will be. This example illustrates that the observed behaviour of a subject that is risk neutral is very different from the observed behaviour of a subject who holds exactly the same beliefs and is risk averse.

We specify a joint likelihood where we simultaneously estimate risk attitudes and subjective beliefs. We assume a CRRA utility function defined by (5.11). The coefficient of risk aversion of this utility function can be estimated using maximum likelihood and a latent structural model of choice, such as Expected Utility Theory (EUT). Suppose that there are K possible outcomes in a lottery. The lotteries we present to our subjects in the risk task of our experiment have $K \leq 3$. The probability for each outcome, p_K , is exogenously induced by the experimenter, so the expected utility of each lottery is the sum of the probability weighted utilities of each outcome in each lottery i :

$$EU_i = \sum_{k=1}^K [p_k \times u_k] \quad (5.12)$$

We calculate the EU for each lottery pair, assuming some trial value for the parameter r , and define the index:

$$\nabla EU = EU_R - EU_L \quad (5.13)$$

where EU_R is the expected utility of the right lottery and the EU_L is the expected utility of the left lottery. This latent index, based on latent preferences, is then linked to observed choices using a cumulative normal distribution function $\Phi(\cdot)$. This probit function takes any argument between $\pm\infty$ and transforms it into a number between 0 and 1. The agent chooses R if $\nabla EU + \varepsilon > 0$, where ε is a normally distributed error term with mean zero and variance σ^2 . Thus we have the probit link function, showing the probability that R is chosen, as

$$\begin{aligned} \text{prob (choose lottery R)} &= P(\nabla EU + \varepsilon > 0) = \\ &P(\varepsilon / \sigma > -\nabla EU / \sigma) = \Phi(\nabla EU / \sigma) \end{aligned} \quad (5.14)$$

The index defined by (3.3) is linked to observed choices by specifying that the R lottery is chosen when $\nabla EU + \varepsilon > 0$, which is implied by (3.4). Thus the likelihood of the observed choices, conditional on EUT and CRRA specifications being true, depends on the estimates of r given the above statistical specification and observed choices. The conditional log likelihood is:

$$\ln L(r; y, \mathbf{X}) = \sum_i [\ln \Phi(\nabla EU) \times I(y_i = 1) + (\ln \Phi(1 - \nabla EU)) \times I(y_i = -1)] \quad (5.15)$$

where $I(\cdot)$ is the indicator function, $y_i=1(-1)$ denotes the choice of the option R (L) lottery in risk aversion task i , and \mathbf{X} is a vector of individual characteristics, reflecting age, sex, and so on, and treatment characteristics, i.e., whether the risk task preceded the belief task or whether the dice in the belief task were presenting in descending or ascending order. The parameter r is defined as a linear function of the characteristics in the vector \mathbf{X} .

Equation 5.15 is the joint density of all the observations in our sample (i.e., the choices made) conditional on r , which we estimate from the data, and the individual and treatment characteristics. An important feature of the model is that it allows subjects to make some errors in their evaluations of the lotteries. The notion that subjects are subject to such errors has already been incorporated in the design in equation (5.14), where we acknowledge that the subjects chooses the lottery with the highest EU with less than certainty. These errors are traditionally interpreted as deriving from sampling error. In addition, we might posit an explicitly behavioural error, where a subject with a particular r may in some cases choose the lottery that provides lower utility simply by confusion. We capture such noise in the data using the Fechner error

specification, used widely in the literature [Hey and Orme (1994)]. This transforms the index in (5.13) to:

$$\nabla EU = \frac{EU_R - EU_L}{\mu} \quad (5.16)$$

where $\mu > 0$ is a structural “noise” parameter used to capture any errors in judgment. As $\mu \rightarrow \infty$, ∇EU in (5.16) converges to 0 for any finite value of EU_R and EU_L , and so the probability of either choice converges to $\frac{1}{2}$. This means that for any given r , the differences between the EU of the two lotteries are less predictive of choices. As $\mu \rightarrow 0$, ∇EU in (5.16) converges to the “undistorted” difference between EU_R and EU_L , so the probability of either choice converges to 0 or 1. This means that for any given r , the differences between the EU of the two lotteries are fully predictive of choices.

Table 5.6 shows the estimates for r and μ when we estimate the risk aversion coefficient, as explained above. We include the dummy variable *task* which equals one if the Risk Attitude (RA) task preceded the Belief Elicitation (BE) task and zero otherwise, and two dummies (*Group2* and *Group3*) to flag the different groups of lottery tasks we presented to our subjects. We also include dummy variables to capture any effect of demographic characteristics⁷¹ on risk attitudes. These are: *female* (equals to 1 if the subject is female and 0 otherwise), *teenager* (1 if subject is less than 20 years old and 0 otherwise), *white* (1 if the subject described himself as white and 0 otherwise), *British* (1 if the subject is a British citizen and 0 otherwise), *high income* (1 if the subjects earn more than £10,000 per year and 0 otherwise), *graduate* (1 if the subject is a postgraduate and 0 otherwise) and *math* (1 if the subject studies Economics, Finance, Engineering, Mathematics, Computer or Physical Sciences and 0 otherwise).

⁷¹ The demographics questionnaire is shown in Appendix 3.

The estimation reveals that on average the subjects in our sample are risk averse. The coefficient of *task* is 0.186 (p -value 0.066), indicating that our subjects become *more* risk averse when the BE task precedes the RA task. These order effects are common in the literature [Harrison et al (2005)] and indicate potential wealth effects. For example, when the RA task follows the BE task subjects already know their payment for the first part of the experiment. This may affect their responses in the RA task. The coefficient of the *Group2* dummy is insignificant, but the coefficient of the *Group3* dummy is -0.114 and significant (p -value 0.062). This means that subjects were less risk averse, if in the RA task they made choices using the third group of lotteries. In terms of demographics, all coefficients are insignificant except *female* (0.199 with p -value 0.007) and *high income* (0.147 with p -value 0.007), showing that females and high income subjects are more risk averse than males and low income subjects.

The analysis in this section shows that the assumption of risk neutrality is not appropriate for our data. This justifies eliciting subjective beliefs by taking into account the curvature of the utility function. In fact, even if we had shown that subjects were statistically risk neutral, the presence of any uncertainty in that claim (viz., a non-zero standard error in the estimate of r) would require that we jointly estimate risk attitudes and beliefs to ensure that we accurately characterize the uncertainty in estimated beliefs.

5.3.3 Eliciting beliefs after controlling for preference towards risk

This section explains in more detail the steps in the procedure to elicit risk attitudes and subjective beliefs jointly. From Table 5.2 a subject who selects to bet on the white box being used for any given bookie receives EU:

$$EU_w = \pi_A \times U(\text{payout if } W \text{ used} | \text{bet on } A) \quad (5.17)$$

where π_A is the subject's belief of the likelihood of event A. The payouts that enter the utility function correspond to the odds given by each particular bookie, as shown by Table 5.2. For example, in Table 5.2, for the first bookie, the individual bets on the white box being used. The EU of this bet is $\pi_A \times U(\pounds 60) + (1 - \pi_A) \times U(\pounds 0)$, where π_A is the subject's personal belief of the white box being chosen initially. Since we observe the choice for each subject and bookie we can calculate the likelihood of that choice given values for r , π_A and μ . We need r to evaluate the utility function in (5.11), we need π_A to calculate the EU of (5.17), and we need μ to calculate the index in (5.16). The joint elicitation problem is to find values for these three parameters that best explain the choices in both the risk *and* the belief tasks. Since probabilities are bounded between [0,1], we constrain the probability estimates to vary in this interval.⁷²

Table 5.7 shows the results when we perform joint estimation. Panel A presents the estimated risk aversion coefficient, and the effects of the dummy variables that flag treatment and subject characteristics. The estimate of the relative risk aversion coefficient is 0.629 (p -value < 0.001), which is larger than the coefficient depicted in Table 5.5 derived from only the responses in the RA task. The coefficient of the dummy *task* is -0.027 (p -value 0.096), which indicates that subjects become less risk averse when the BE task precedes the RA task. The dummy for high income subjects is positive (0.033) and significant (p -value 0.040), showing that subjects from richer families are slightly more risk averse. All other demographic-related dummies do not have any significant effect on risk attitudes.

⁷² We use the technique discussed in Harrison and Rutström (2008) and transform the elicited parameter, say k , as $\pi_A = 1/(1 + e^k)$. This transformation allows the estimated parameter, k , to vary freely between $\pm\infty$, and returns a value in the closed interval [0, 1].

In terms of elicited beliefs, the dummies *task* and *Descending* are negative (-0.025 and -0.026 respectively) and significant (p -values of 0.037 and 0.026 respectively), showing that that elicited probabilities are lower when the BE task precedes the RA task or when the samples in the BE task are presented in descending order.

Panel B of Table 5.7 shows the elicited probabilities for the various patterns, their strength- weight characteristics, and the “correct” posterior probabilities implied by applying Bayes Rule. The various dice patterns are placed into 6 groups according to their correct probabilities. If subjects use Bayes Rule to update their beliefs, elicited probabilities in each of the 6 groups of patterns should equal this correct probability. If, however, subjects update their beliefs according to the hypothesis advanced by GT, we should observe that in each group elicited beliefs are updated more strongly as the strength of the pattern increases, and that subjects are overconfident (underconfident) when strength is high (low) and weight is low (high).

The hypothesis that elicited probabilities are equal to the probabilities implied by applying Bayes rule is rejected for all patterns. This means that regardless of the strength or the weight of the particular pattern, elicited beliefs *are not* updated according to Bayes Rule. The GT hypothesis predicts that within groups of patterns that yield the same posterior, elicited beliefs will be higher when the strength of the pattern is high. We test this prediction by comparing elicited probabilities in the 6 groups of patterns that have the same posterior, but differ in strength and weight. The hypothesis that elicited probabilities in the six groups are equal is rejected in all cases. The results show that elicited probabilities in each group increase, as the strength of the pattern increases. For example, consider the patterns that yield a posterior probability of 0.88. The pattern with the lowest strength is (11, 6), followed by the pattern (7, 2) and the pattern (5, 0). Elicited probabilities for these patterns are 0.569, 0.633 and 0.655 respectively, showing a significant difference of 8.6 percent between elicited probabilities for the

low/high strength patterns. The same relationship between strength and elicited probabilities can be seen from the remainder groups of patterns in Table 5.6. For example, for patterns that yield a posterior probability of 0.77, the elicited probability for the high strength pattern (3,0) is 4.7 percentage points higher when compared to the low strength pattern (10,7). This finding suggests that subjects, contrary to Bays Rule, update their beliefs more strongly as the strength of the pattern increases, as suggested by GT.

5.3.4 Estimating maximum likelihood parameters of the effect of strength and weight on elicited beliefs

We now turn to estimating the coefficients of α and β , as shown in Equation 5.8. We can estimate these coefficients in a structural maximum likelihood model of betting behaviour. Assume, for now, that $w > b$. Then if we know α and β we can *evaluate* the RHS of (5.8) as γ , and obtain:

$$\log \left[\frac{\text{Log}\left(\frac{P(W/D)}{P(B/D)}\right)}{\text{Log}\left(\frac{0.6}{0.4}\right)} \right] = \gamma \quad (5.18)$$

Therefore:

$$\frac{\text{Log}\left(\frac{P(W/D)}{P(B/D)}\right)}{\text{Log}\left(\frac{0.6}{0.4}\right)} = e^\gamma \quad (5.19)$$

Rearranging:

$$\text{Log}\left(\frac{P(W/D)}{P(B/D)}\right) = e^\gamma \text{Log}\left(\frac{0.6}{0.4}\right) \quad (5.20)$$

Transforming both sides:

$$\frac{P(W/D)}{P(B/D)} = e^{e^\gamma \log\left(\frac{0.6}{0.4}\right)} \quad (5.21)$$

Let the RHS of (5.21) be δ , so we have:

$$\frac{P(W/D)}{P(B/D)} = \delta \quad (5.22)$$

Which immediately gives us:

$$P(W/D) = \frac{\delta}{1 + \delta} \quad (5.23)$$

This instead of estimating $P(W/D)$ directly, we can estimate α and β , and given the characteristics of the choice (n , w , b , and the prior 0.6) calculate $P(W/D)$ and evaluate the EU of the bet as before.

One point when we do the above estimations is that the specifications in (5.18-5.23) are fine when the white dice exceed the blue dice. However, when the blue dice exceed the white dice, $P(B/D) > P(W/D)$. This means that the ratio of these probabilities in the numerator of the left hand side expression of (5.8) will be less than one. Being less than one, we can only take its logarithm once, as logs of negative numbers are undefined. Similarly, when blue $>$ white, the strength expression on the right hand side of (5.8) will be negative, so its logarithm is also undefined.

In order to get around this issue we always define the log likelihood ratio in the numerator of the left hand side expression in (5.8) with the probability of the *most* likely box, as suggested by the sample revealed, being the numerator. Therefore, when we have more blue dice than white dice, in which case the blue box is more probable, we transform (5.8) to:

$$\text{Log} \left[\frac{\text{Log} \left(\frac{P(B/D)}{P(W/D)} \right)}{\text{Log} \left(\frac{0.6}{0.4} \right)} \right] = a \log(n) + \beta \log \left(\frac{(b-w)}{n} \right) \quad (5.24)$$

In GT's design, because their data were hypothetical, they always supported the same hypothesis, so such issues did not arise. However, because our data are randomly generated in the experiment, we need to define the variables in (5.8) according to the particular pattern that emerges, so we can estimate a and b .

If we use (5.24) and make the same transformations as before we get:

$$P(B/D) = \frac{1 + \delta}{\delta} \quad (5.25)$$

Table 5.8 shows the maximum likelihood estimates of alpha and beta, the coefficient of risk aversion and the behavioural error term. The coefficients of both strength and weight are significantly different from 0. More importantly, they are significantly different from each other in the direction predicted by the GR hypothesis, i.e., that $\alpha < \beta$ ($\alpha = 0.267$ and $\beta = 0.738$, both with p -value < 0.001). This result provides further support to the hypothesis advanced by GT that

subjective probabilities are more responsive to the strength of the information presented than the weight.

Our analysis can be seen to extend that of GT's as we estimate the parameters in a model of maximum likelihood, whereas they use their elicited beliefs as the dependent variable in (5.8) and estimate α and β using OLS. Econometrically this can be seen as problematic because elicited probabilities are essentially estimates of latent beliefs and as such they entail uncertainty. Using them as dependent variables in a regression does not account for this uncertainty, which may lead to false conclusions in terms of the relevant coefficients. Our approach remedies this limitation as we estimate α and β from the choices made in a more rigorous manner using maximum likelihood. In addition we allow for risk aversion, thus controlling for the effect of the concavity of the utility function on the effects of strength and weight on elicited beliefs. The fact that the GT hypothesis survives these controls demonstrates its robustness as an alternative decision making model to Bayesian updating.

5.4. Further Discussion

Our study belongs in a category of studies in experimental economics that attempt to examine whether the various behavioural biases documented by psychologists are robust. This is an important exercise, especially with the expansion of the field of behavioural finance, which incorporates these biases into asset pricing theories. If these biases are not genuine, and reflect the shortcomings and limitations of the experimental design employed by psychologists, then their application in economics will produce erroneous results.

This is not to say that the results produced by experiments in psychology are not interesting to economists. The issue is that psychologists' primary interest is to analyse peoples'

first response to a situation, and not their *economic behaviour*, which arguably involves careful deliberation and real consequences. Therefore, in order to apply the results from psychology into economics, it is important to establish their validity in an environment that simulates an economic decision.

Our finding that the bias in subjects' beliefs is much smaller compared to the bias documented in the original study of Griffin and Tversky (1992) suggests that the methodological design of an experiment can affect the outcome. However, at the same time, the fact that the GT hypothesis survives the scrutiny that is usually advocated by economists indicates that psychological studies produce robust results. In any case the general message of our analysis is that psychological theories should not be automatically dismissed from economics based on methodological limitations, but rather should stimulate further research.⁷³

5.5. What do these results mean for financial markets?

This chapter has highlighted that beliefs are more sensitive to the extremity of the information presented rather than its predictive validity. This section explores the relationship of this finding with some of the empirical finance literature.

Tetlock (2006) shows that the power of the language used in a very popular column in the Wall street journal predicts market wide movements, indicative of overreaction. Similarly, Barber and Loeffler (1993) show that stocks that are recommended by a panel of analysts in the 'dartboard column' exhibit statistically and economically significant abnormal returns, and subsequent reversals. This evidence suggests that investors focus too much on eye-grabbing information signals without too much consideration on their predictive validity, as we have

⁷³ The role of these methodological limitations is vastly debated between psychologists and economists, with mixed results. For example, Read (2005) suggests that monetary incentives do not guarantee that subjects will expense a higher effort in the experimental task.

documented in the experiment. Shiller (2004) suggests that the media is a source of high strength and low weight information, and acts as a propagating factor of mispricings during stock market booms and crashes. For example, stories of successful investments in the late 90's exacerbated the market's potential for profits and contributed to the *dot com* boom and subsequent crash. A striking example of the irrationality of the expectations that prevailed at the time is given by Cooper et al. (2001) who show that companies that simply changed their name to imply a technology involved firm during the dot com bubble, experienced higher abnormal returns. This reflected the consensus which was that due to the rapid of expansion of the internet, the profit capacity of technological firms was limitless.

Odean and Barber (2005) also support the salient information hypothesis. They use a data sample comprised of individual transactions that allow analysis of the buying and selling behaviour of investors. They motivate their study by acknowledging that the buying decision is fundamentally different than the selling decision. Individual investors when considering stocks to buy are faced with thousands of alternatives and are thus unable to consider all of them. Thus they rely on attention that is triggered by the salience of news. However when selling, people only consider the stocks which they already own, a much easier task. They use proxies for salient news, for example extreme trading volume and extreme returns, and whether the stock appeared in the news (using the Nexis database). They find that buying behaviour is strongly affected by salient news, which would score highly on the strength dimension in the terms of Griffin and Tversky (1992).⁷⁴

Klibanoff, Lamont and Wizman (1998) test the performance of U.S closed-end country funds. In their sample they include funds, which entail a commonly traded stock, that invest in

⁷⁴ The behaviour of professional investors was not found to be prone to this kind of 'attention' bias.

39 countries. A closed end funds entails a Net Asset Value. The authors test how the market price of the stocks included in the fund reacts to changes to the NAV, when salient news is present and otherwise. In a 'rational' context the elasticity of the asset price and the NAV should be 1 (or close to it) with or without news, as the latter is a measure of the fundamental value of the firm. Unconditionally the authors observe that the elasticity is 0, 64 that indicates that the price is not very responsive to NAV. However in the presence of salient information (appeared in the cover of Times) for a particular country that potentially affects the NAV, the price elasticity of NAV for the particular fund rises significantly to 83%. This shows that the response of the market, in the presence of salient news is much more dramatic.

Fehle, Tsyplakov and Zdorovtsov (2005) document that companies that advertise during the Super Bowl in the U.S (an event with an immense audience), experience a positive effect on their price consequently. This is a clear irrationality because advertising is a piece of information that glamorizes a company without yielding any substantive information in terms of earnings capacity. Rather the existence of a salient yet uninformative event appears to affect prices.

All these results suggest that the characteristics of the information set available influence the resulting choices, much like we have documented in our experiment. Investors focus excessively on the saliency of the information, and tend to undermine its predictive validity.

5.5 Conclusion

Bayesian updating is at the core of rational decision making. A large number of studies in psychology and decision making show evidence that people systematically violate this paradigm [see Kahneman, Slovic and Tversky (1982)]. This evidence has been heavily applied in the area of behavioural finance. However, supporters of neoclassical theories often view these results

sceptically because the psychological experiments that document them generally do not account for certain features that have been shown to affect subjects' responses, such as financial incentives, naturally occurring information and risk attitudes. Given that these factors have been shown to affect subjects' economic behaviour, it is vital that we examine whether these behavioural biases are robust to these controls before we apply them in economic models of asset prices.

Griffin and Tversky (1992) present evidence to support one such hypothesis of non-Bayesian updating. This hypothesis suggests that when the information used is high extremity and low predictive validity subjects overreact to it, whereas if it is of low extremity and high predictive validity subjects underreact to it. Financial theories (Barberis et al 1998, Sorescu and Subrahmanyam 2006) have utilised the capacity of this hypothesis to identify the conditions that spur over and under reaction amongst investors to explain the "anomalous" behaviour of prices.

However, as mentioned, economists view this evidence sceptically because the studies that document them entail certain methodological shortcomings. Particularly the GT study can be criticised on three accounts:

1. The payments of subjects were not clearly linked to their performance, so subjects were not incentivised to provide accurate responses.
2. The information used by subjects to update their beliefs was hypothetical.
3. Subjects were assumed to be risk neutral throughout

In this study we contribute to the literature by taking a step back and re-examining whether the GT hypothesis holds after we address all these limitations. In this manner we

provide evidence whether this theory provides a viable model of subjective probability formation, which can be used to explain the behaviour of asset prices.


We clearly link the performance of subjects to their payments, so that they are incentivised to perform the task to the best of their ability. Secondly, the information that subjects used to update their beliefs is endogenously generated in the experiment so that subjects are not required to visualize or imagine anything. Lastly, we go beyond risk neutrality and estimate subjective beliefs whilst allowing for risk aversion.

Our results broadly support the hypothesis of GT. We find that elicited probabilities are higher when the information used by subjects is of high strength and low weight. Particularly, we find that the effect of strength on elicited probabilities outweighs the effect of weight by roughly a factor of 2. This result is a clear violation of Bayesian updating as this rule prescribes that these two elements of the information set should have an *equal* effect on elicited probabilities. This evidence suggests that the evidence provided in the original study of GT were not due to the methodological shortcomings of their experimental design, but rather reflect a genuine feature of human behaviour under uncertainty.

However, our design does present some key differences. Firstly, we show that the magnitude of the bias is reduced substantially in our study. For example in Table 5.1 of their paper they find that for samples with posterior probability of 0.77, overconfidence is equal to 8% for the high strength low weight sample (3,0) and underconfidence is 17.5% for the low strength high weight (10,7). This suggests a bias in the magnitude of 25.5%. In our analysis, under the assumption of risk neutrality the corresponding figure is 11.4%. Once we allow for risk aversion, as theory and evidence prescribe, the bias reduces further to 4.7% (similar results hold for all patterns). These differences are striking, and suggest that the implementation of an experimental design more compatible with economic decision making reduces the strength-weight heuristic


found by GT on elicited beliefs. This demonstrates that it is prudent to confirm whether behavioural biases arise in experiments that successfully simulate economic decisions before we actually use them in asset pricing models.

Table 5.1: An example of a lottery pair



13% chance of £0 (numbers 1-13)
87% chance of £5 (numbers 14-100)

Your choice:



37% chance of £0 (numbers 1-37)
37% chance of £5 (numbers 38-74)
26% chance of £15 (numbers 75-100)

Figure 5.1 Boxes and Ten-Sided Dice.



Table 5.2: The table in which subjects place their bets

Bookie	Stake	Odds offered		Earnings including the stake of £3		I will bet on (circle)
		White	Blue	White	Blue	
1	£3	20.00	1.05	£60.00	£3.15	<u>W</u> B
2	£3	10.00	1.11	£30.00	£3.33	<u>W</u> B
3	£3	6.67	1.18	£20.00	£3.54	<u>W</u> B
4	£3	5.00	1.25	£15.00	£3.75	<u>W</u> B
5	£3	4.00	1.33	£12.00	£4.00	W <u>B</u>
6	£3	3.33	1.43	£10.00	£4.29	W <u>B</u>
7	£3	2.86	1.54	£8.58	£4.62	W <u>B</u>
8	£3	2.50	1.67	£7.50	£5.00	W <u>B</u>
9	£3	2.22	1.82	£6.66	£5.46	W <u>B</u>
10	£3	2.00	2.00	£6.00	£6.00	W <u>B</u>
11	£3	1.82	2.22	£5.46	£6.66	W <u>B</u>
12	£3	1.67	2.50	£5.00	£7.50	W <u>B</u>
13	£3	1.54	2.86	£4.62	£8.58	W <u>B</u>
14	£3	1.43	3.33	£4.29	£10.00	W <u>B</u>
15	£3	1.33	4.00	£4.00	£12.00	W <u>B</u>
16	£3	1.25	5.00	£3.75	£15.00	W <u>B</u>
17	£3	1.18	6.67	£3.54	£20.00	W <u>B</u>
18	£3	1.11	10.00	£3.33	£30.00	W <u>B</u>
19	£3	1.05	20.00	£3.15	£60.00	W <u>B</u>

Table 5.3: Composition, strength and weight characteristics, Bayesian probability and number of observations of the various dice patterns

Total Dice (<i>n</i>)	No White (<i>w</i>)	No blue (<i>b</i>)	Strength $/(w-b)/n/$	Bayesian Probability	No of observations
5	0	5	1.00	0.12	72
9	2	7	0.56	0.12	110
17	6	11	0.29	0.12	47
3	0	3	1.00	0.23	77
5	1	4	0.60	0.23	218
9	3	6	0.33	0.23	112
17	7	10	0.18	0.23	106
3	1	2	0.33	0.4	110
5	2	3	0.20	0.4	412
9	4	5	0.11	0.4	117
17	8	9	0.06	0.4	82
3	2	1	0.33	0.6	154
5	3	2	0.20	0.6	394
9	5	4	0.11	0.6	87
17	9	8	0.06	0.6	44
3	3	0	1.00	0.77	55
5	4	1	0.60	0.77	227
9	6	3	0.33	0.77	95
17	10	7	0.18	0.77	73
5	5	0	1.00	0.88	45
9	7	2	0.56	0.88	53
17	11	6	0.29	0.88	81

Figure 5.2: Distribution of risk neutral beliefs from raw data for each pattern for each group with same posterior probability

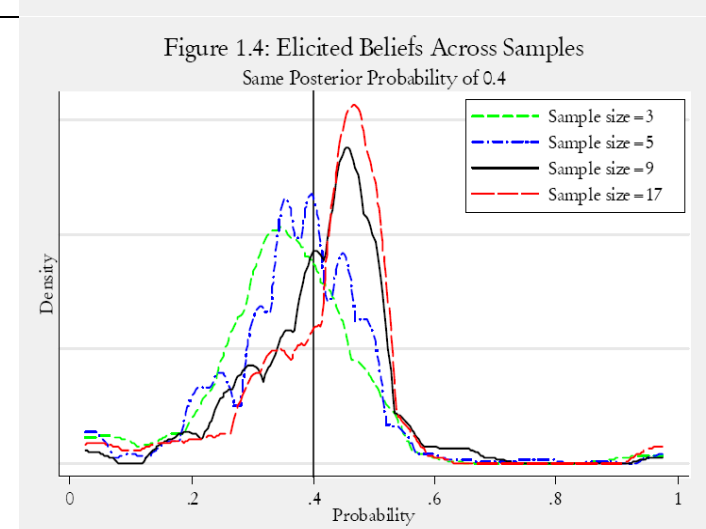







Table 5.4: Maximum Likelihood estimates of Risk Neutral Subjective Beliefs

Interval regression estimates of elicited probabilities

Panel A: Effect of task order and ascending/descending designs on elicited probabilities.

Parameter	Estimate	St. error	p-value	95% conf. interval	
π_A					
task	0.0147	0.007	0.043	0.000	0.029
Descending	0.0186	0.007	0.009	0.005	0.033

Panel B: Strength, weight, Bayesian probability and elicited probabilities

White dice	Blue dice	Strength $/(w-b)/n/$	Weight (n)	Bayesian Prob.	Elicited Prob.	St. error	95% conf. interval	
0	5	1.00	5	0.12	0.122	0.027	0.087	0.156
2	7	0.56	9	0.12	0.164	0.026	0.138	0.189
6	11	0.29	17	0.12	0.250	0.032	0.198	0.301
0	3	1.00	3	0.23	0.201	0.029	0.170	0.232
1	4	0.60	5	0.23	0.222	0.024	0.198	0.246
3	6	0.33	9	0.23	0.248	0.026	0.218	0.278
7	10	0.18	17	0.23	0.320	0.027	0.291	0.348
1	2	0.33	3	0.40	0.333	0.025	0.304	0.361
2	3	0.20	5	0.40	0.352	0.021	0.330	0.373
4	5	0.11	9	0.40	0.393	0.022	0.364	0.421
8	9	0.06	17	0.40	0.406	0.023	0.373	0.439
2	1	0.33	3	0.60	0.623	0.020	0.592	0.653
3	2	0.20	5	0.60	0.586	0.020	0.562	0.609
5	4	0.11	9	0.60	0.551	0.023	0.518	0.585
9	8	0.06	17	0.60	0.552	0.028	0.506	0.598
3	0	1.00	3	0.77	0.773	0.019	0.734	0.812
4	1	0.60	5	0.77	0.714	0.021	0.676	0.752
6	3	0.33	9	0.77	0.673	0.032	0.618	0.727
10	7	0.18	17	0.77	0.659	0.027	0.616	0.700
5	0	1.00	5	0.88	0.805	0.029	0.749	0.861
7	2	0.56	9	0.88	0.793	0.026	0.744	0.842
11	6	0.29	17	0.88	0.716	0.026	0.664	0.767

Table 5.5: Comparison of bets between a risk neutral and a risk averse agent with the same beliefs.

Bookie	B		EV(W)	EV(B)	BET RN	EU(W)	EU(B)	BET RA
	Payout of W	Payout of B						
1	60	Not 60	45	0.79	W	10.14	0.95	W
2	30	3.33	22.5	0.83	W	7.53	0.98	W
3	20	3.54	15	0.89	W	6.32	1	W
4	15	3.75	11.25	0.94	W	5.59	1.03	W
5	12	4	9	1	W	5.08	1.06	W
6	10	4.29	7.5	1.07	W	4.69	1.09	W
7	8.58	4.62	6.44	1.16	W	4.4	1.12	W
8	7.5	5	5.63	1.25	W	4.15	1.16	W
9	6.66	5.46	5	1.37	W	3.94	1.21	W
10	6	6	4.5	1.5	W	3.77	1.26	W
11	5.46	6.66	4.1	1.67	W	3.62	1.31	W
12	5	7.5	3.75	1.88	W	3.48	1.38	W
13	4.62	8.58	3.47	2.15	W	3.37	1.47	W
14	4.29	10	3.22	2.5	W	3.26	1.56	W
15	4	12	3	3	W	3.17	1.69	W
16	3.75	15	2.81	3.75	B	3.08	1.86	W
17	3.54	20	2.66	5	B	3	2.11	W
18	3.33	30	2.5	7.5	B	2.93	2.51	W
19	3.15	60	2.36	15	B	2.86	3.38	B

Table 5.6: Maximum Likelihood estimates of risk aversion

Assuming CRRA and EUT preferences

Parameter	Estimate	Standard error	p-value	95% conf. interval	
constant	0.458	0.110	0.00	0.243	0.674
Group2	0.015	0.057	0.79	-0.096	0.126
Group3	-0.114	0.062	0.066	-0.236	0.007
Raorder	0.186	0.062	0.002	0.065	0.307
Female	0.199	0.074	0.007	0.055	0.344
Teenager	-0.078	0.061	0.2	-0.199	0.042
White	0.169	0.117	0.149	-0.061	0.399
British	-0.227	0.138	0.101	-0.498	0.044
High Income	0.147	0.054	0.007	0.041	0.254
Graduate	-0.111	0.103	0.279	-0.313	0.090
Math	-0.106	0.070	0.134	-0.244	0.032
μ_{RA}	0.901	0.08	0.000	0.745	1.06

Table 5.7: Maximum Likelihood estimates of Risk Averse Subjective Beliefs
Assuming EUT and CRRA utility function.

Panel A: Risk aversion coefficients, demographics and elicited beliefs

Parameter	Estimate	St. error	p-value	95% conf. interval	
<i>r</i>	0.629	0.066	0	0.5	0.757
task	-0.027	0.016	0.096	-0.059	0.005
Female	0.021	0.017	0.228	-0.013	0.055
Teenager	-0.011	0.017	0.521	-0.044	0.022
White	0.026	0.029	0.371	-0.031	0.084
British	0.053	0.062	0.388	-0.068	0.174
High Income	0.033	0.016	0.04	0.001	0.065
Graduate	-0.006	0.037	0.879	-0.078	0.067
Math	-0.002	0.017	0.912	-0.036	0.032
π_A					
task	-0.025	0.012	0.037	-0.048	-0.001
Descending	-0.026	0.012	0.026	-0.049	-0.003
Female	0.002	0.016	0.925	-0.031	0.034
Teenager	0.01	0.011	0.336	-0.011	0.032
White	-0.018	0.025	0.471	-0.067	0.031
British	-0.017	0.032	0.601	-0.079	0.046
High Income	-0.016	0.014	0.243	-0.044	0.011
Graduate	0.028	0.025	0.269	-0.021	0.077
Math	0.003	0.015	0.858	-0.026	0.031

Table 5.7: Continued

Panel B: Pattern strength, weight and elicited risk averse beliefs.

No white	No Blue	Strength	Weight	Bayesian	Elicited Prob.	St. error	95% confidence interval	
(w)	(b)	$ (w-b)/n $	(n)	Prob.	(π_A)			
2	1	0.33	3	0.6	0.522	0.009	0.505	0.54
3	2	0.2	5	0.6	0.506	0.007	0.493	0.52
5	4	0.11	9	0.6	0.491	0.009	0.474	0.508
9	8	0.06	17	0.6	0.499	0.01	0.48	0.519
3	0	1	3	0.77	0.606	0.013	0.58	0.632
4	1	0.6	5	0.77	0.57	0.011	0.549	0.59
6	3	0.33	9	0.77	0.552	0.015	0.522	0.581
10	7	0.18	17	0.77	0.542	0.01	0.522	0.562
5	0	1	5	0.88	0.655	0.024	0.607	0.703
7	2	0.56	9	0.88	0.633	0.018	0.597	0.669
11	6	0.29	17	0.88	0.569	0.015	0.54	0.599
1	2	0.33	3	0.4	0.425	0.007	0.412	0.438
2	3	0.2	5	0.4	0.43	0.006	0.419	0.441
4	5	0.11	9	0.4	0.443	0.007	0.43	0.455
8	9	0.06	17	0.4	0.442	0.007	0.429	0.456
0	3	1	3	0.23	0.38	0.007	0.367	0.394
1	4	0.6	5	0.23	0.389	0.005	0.379	0.4
3	6	0.33	9	0.23	0.398	0.007	0.385	0.411
7	10	0.18	17	0.23	0.421	0.007	0.408	0.434
0	5	1	5	0.12	0.359	0.008	0.343	0.374
2	7	0.56	9	0.12	0.374	0.007	0.361	0.388
6	11	0.29	17	0.12	0.394	0.008	0.379	0.41
	5	1	5	0.12	0.359	0.008	0.343	0.374

Table 5.8: Maximum Likelihood estimates of coefficients of strength, weight and risk aversion

Assuming EUT and CRRA preferences

	Estimate	Standard error	p value	95% conf. interval	
r	0.699	0.015	<0.001	0.669	0.729
alpha	0.267	0.033	<0.001	0.201	0.332
beta	0.738	0.038	<0.001	0.662	0.814
Lnmu	-1.262	0.124	<0.001	-1.51	-1.022

Appendix 5.1: Instructions for the Risk Aversion task

Stage 2: INSTRUCTIONS

We will now continue with Stage 2 of the experiment.

This stage is about choosing between lotteries with varying prizes and chances of winning. You will be shown 20 lottery pairs, and from each pair you will choose the lottery you prefer. You will actually get the chance to play one of the lotteries you choose, and will be paid according to the outcome of that lottery, so you should think carefully about your preferences.

On the accompanying sheet there is an example lottery pair.

The outcome of the lotteries will be determined by the roll of a 100-sided die that is numbered from 1 to 100. The numbers that will determine each outcome are shown below each lottery.

In the example the left lottery pays five pounds (£5) if the number is between 1 and 40 (a 40% chance), and it pays fifteen pounds (£15) if it is between 41 and 100 (a 60% chance).

The lottery on the right pays five pounds (£5) if the number drawn is between 1 and 50 (a 50% chance), ten pounds (£10) if the number is between 51 and 90 (a 40% chance), and fifteen pounds (£15) if the number is between 91 and 100 (a 10% chance).

The size of the pie slices represent the chances of earning each payoff.

Each lottery pair will be shown on a separate sheet of paper. On each sheet you should indicate your preferred lottery by ticking the appropriate box. After you have worked through all the lottery pairs, please raise your hand.

You will then roll a 20-sided die to determine which pair of lotteries will be played out, and the 100-sided die to determine the outcome of the chosen lottery.

For instance, suppose the lottery on the accompanying page was chosen to pay off and you rolled a 42 on the 100-sided die. If you had picked the lottery on the left you would win £15, while if you had picked the lottery on the right you would have won £5.

Therefore, your payoff is determined by three things:

- which lottery pair is chosen to be played out using the 20-sided die;
- which lottery you selected, the left or the right, for the chosen lottery pair;
- the outcome of that lottery when you roll the 100-sided die.

This is not a test of whether you can pick the best lottery in each pair, because none of the lotteries are necessarily better than the others. Which lotteries you prefer is a matter of personal taste.

Please work silently, and think carefully about each choice.

All payoffs are in cash, and are in addition to the £5 show-up fee that you receive just for being here.

Appendix 5.2: Instructions for the Belief Elicitation task.

Stage 3: INSTRUCTIONS

What you will do

In this stage of the experiment you will be betting on the outcomes of uncertain events. Usually we bet on events like football matches or elections, but in this task the events will be random choices made by the experimenter between two boxes, one blue and the other white. The experimenter will not tell you which box was chosen. At the start each box will have the same chance of being chosen, but once it has been chosen the experimenter will give you some information to help you work out the chances that it was blue or white. Armed with this information, you will make bets on which box was chosen.

The procedure, which is summarized on the accompanying picture, is as follows. The experimenter will first choose the box by rolling a 6-sided die with three blue and three white sides. If blue comes up he will choose the blue box, if white comes up he will choose the white one.

Both the white and blue boxes contain several dice, each having 10 sides. Both boxes have the same number of dice, which will vary over the course of the experiment. The dice in the blue box always have 6 blue sides and 4 white ones, while those in the white box have 4 blue sides and 6 white ones.

The experimenter will roll all the dice in the chosen box and tell you how many blue and white sides came up. He will not tell you which box was chosen.

Because the dice in the blue box have more blue sides than those in the white box, knowing the number of blue and white sides that come up can help you work out the chances that each box was chosen. For example, if more blue sides come up this means it is more likely to be the blue box, and if more white sides come up it is more likely to be the white box.

Once you have the information about the dice rolls, you will then make bets on which box was chosen.

About betting

You will be making bets with several betting houses or “bookies,” just as you might bet on a football game or a horse race.

To familiarize you with betting, we will illustrate how it works with the example of a horse race.

Imagine a two horse race between Blue Bird and White Heat. Several bookies offer different odds for both horses. The table below shows the odds offered by three bookies along with the amounts they would pay if you staked £10 on the *winning* horse. The earnings are calculated by multiplying the odds by the stake. In this experiment you will be making bets on which box was chosen using a table like this. **At this point you should take some time to study the table.**

Bookie	Stake	Odds offered		Earnings including the stake of £10	
		Blue Bird	White Heat	Blue Bird	White Heat
A	£10	5.00	1.25	£50.00	£12.50
B	£10	3.33	1.43	£33.33	£14.30
C	£10	2.00	2.00	£20.00	£20.00

Below are three important points about betting.

1. **Your belief about the chances of each outcome is a personal judgment that depends on information you have about the different events.** For the horse race, you may have seen previous races or read articles about them. In the experiment the information you have about whether the blue or white box was chosen will be how many blue and white faces came up.
2. **Even if you believe Event X is more likely to occur than Event Y, you may want to bet on Y because you find the odds attractive.** For example, even if you believe White Heat is most likely to win you may want to bet on Blue Bird because you find the odds attractive. To illustrate, suppose you personally believe that Blue Bird has a 40% chance of winning and White Heat has a 60% chance of winning. This means that if you bet £10 on Blue Bird with Bookie A you believe there is a 40% chance of receiving £50.00 and a 60% chance of receiving nothing. You may find this more attractive than betting on White Heat, which you believe offers a 60% chance of 12.50 and a 40% chance of nothing.
3. **Your choices might also depend on your willingness to take risks or to gamble.** There is no right choice for everyone. In a horse race you might want to

bet on the longshot since it will bring you more money if it wins, but you also might want to bet on the favorite since it is more likely to win something.

For each bookie, whether you would choose to bet on Blue Bird or White Heat will depend on three things: your judgment about how likely it is each horse will win, the odds offered by the bookie, and how much you like to gamble or take risks.

Your choices

Now you are familiarized with odds, we can go back to the experimental betting task. Recall that the experimenter will first make a random choice of a blue or white box. Then he will roll the dice in the chosen box and tell you how many white and blue sides came up. Then you will consider the chances that the box chosen was blue or white, and make a series of bets.

You have a booklet of record sheets. Each record sheet shows the bookies you will be dealing with, and the odds they offer. There are 19 bookies on each sheet, and each offer different odds for the two outcomes. **Take a minute to look at one such record sheet, shown on the next page.**

There will be 30 separate events, and 19 bookies offer odds for each event. **You will make bets at all 19 bookies for all 30 events.**


For each bet, you have a £3 stake, and the record sheet shows the payoffs you will receive if you bet on the box that was actually chosen.

There is a separate record sheet for each of the 30 events. On each sheet you should circle W or B to indicate the bet you want to make with **all 19 bookies.**

One and only one of the bets in the entire experiment will pay off for real. Therefore, please consider each bet as if it is the only one that will be paid out. After you have placed all your bets, you will roll a 30-sided die to determine which event will be played out, and a 20-sided die to determine which bookie will determine your earnings.

All payoffs are in cash, and are in addition to the £5 show-up fee that you receive just for being here.

A diagrammatical illustration of the betting task provided to subjects.



Appendix 5.3. The Demographics questionnaire completed by subjects.

In this survey most of the questions asked are descriptive. We will not be grading your answers and your responses are completely confidential. Please think carefully about each question and give your best answers.

1. What is your age? _____ years
2. What is your sex? (Circle one number.)
01 Male 02 Female
3. Which of the following categories best describes you? (Circle one number.)
01 British
02 Irish
03 Any other white background
04 White and Black Caribbean
05 White and Black African
06 White and Asian
07 Any other mixed background

08 Indian
09 Pakistani
10 Bangladeshi
11 Any other Asian background
12 Caribbean
13 African
14 Any other Black background
15 Chinese
16 Any other ethnic group
17 Not stated
18 Prefer not to say
4. What is your main field of study? (Circle one number.)
01 Accounting
02 Economics
03 Finance
04 Business Administration, other than Accounting, Economics, or Finance
05 Education
06 Engineering
07 Health and Medicine
08 Biological and Biomedical Sciences
09 Math, Computer Sciences, or Physical Sciences
10 Social Sciences or History
11 Law
12 Psychology
13 Modern Languages and Cultures
14 Other Fields
5. What is your year of studies? (Circle one number.)
01 First year 04 Masters

02 Second year
03 Third year

05 Doctoral

6. What is the **highest** level of education you expect to **complete**? (Circle one number)

01 Bachelor's degree
02 Master's degree
03 Doctoral degree
04 Professional qualification

7. As a percentage, what is your current average mark if you are doing a Bachelor's degree, or what was it when you did a Bachelor's degree? This mark should refer to all your years of study for this degree, not just the current year. Please pick one by rounding up or down to the nearest number:

01 Above 70%
02 Between 60 - 69%
03 Between 50 - 59%
04 Between 40 - 49%
05 Less than 40%
06 Have not taken courses for which grades are given.

8. What is your citizenship status?

01 British Citizen
02 EU Citizen (non-British Citizen)
03 Non-EU Citizen

9. Are you currently:

01 Single and never married?
02 Married?
03 Separated, divorced or widowed?

10. How many people live in your household? Include yourself, your spouse and any dependents. Do not include your parents or roommates unless you claim them as dependents.

11. Please circle the category below that describes the total amount of income before tax earned in the calendar year 2007 by the people in your household (as "household" is defined in question 10).

[Consider all forms of income, including salaries, tips, interest and dividend payments, scholarship support, student loans, parental support, social security, alimony, and child support, and others.]

01 Less than £10,000
02 £10,000 – £19,999
03 £20,000 – £29,999
04 £30,000 - £49,999

05 Over £50,000

12. Please circle the category below that describes the total amount of income before tax earned in the calendar year 2007 by your parents.
[Consider all forms of income, including salaries, tips, interest and dividend payments, social security, alimony, and child support, and others.]

01 Less than £10,000
02 £10,000 – £19,999
03 £20,000 – £29,999
04 £30,000 - £49,999
05 Over £50,000
06 Don't Know

13. Do you currently smoke cigarettes? (Circle one number.)

00 No
01 Yes

If yes, approximately how much do you smoke in one day? _____ cigarettes.

6. Conclusion

Behavioural finance is a relatively new field that attempts to apply concepts from psychology and behavioural economics in financial markets. By doing so it has been able to explain phenomena that appear puzzling to the traditional neoclassical model. In this thesis we expanded the literature in behavioural finance, by drawing from findings in decision science and proposing behavioural hypotheses that examine how investors use information and invest in financial assets.

In the third chapter, we examine whether ambiguity aversion affects the pricing of analyst forecasts. The importance of analyst earnings forecasts on investment decision is highlighted by Dreman (1998) who states that *“Investors, from the managers of multi-billion-dollar pension funds to the average Joe, act quickly on changes in analyst forecasts, which are flashed immediately by the financial media. Near-term earnings estimates, as noted, are the major trigger for investment decisions today”*, (Dreman (1998), pp 90). Given the large effects that analyst forecasts exert on investment decisions and asset prices Kothari (2001) states that *“It is important to develop refutable hypotheses on the basis of behavioural theories of inefficient financial markets and to perform tests [in terms of the efficiency with which analyst forecasts are transmitted into prices] that discriminate between efficient and inefficient market hypotheses”*, Kothari (2001, pp120-121). Responding to this observation we examine whether ambiguity aversion affects the pricing of analyst forecasts, inducing price predictability.

Measuring ambiguity using company size, we find that prices around analyst forecasts suggest the presence of ambiguity aversion. When investors cannot confidently estimate the behaviour of forecast accuracy (something we illustrate quantitatively), they respond pessimistically, setting prices too low and inducing an ambiguity premium. This study expands

the analyst forecast literature, by examining a behavioural hypothesis of the pricing of analyst forecasts. In addition, it explains the size premium, an important CAPM anomaly. Finally, it is first empirical study of ambiguity aversion in the literature.

In the fourth chapter we examine whether investors sentiment, measured using the survey of consumer confidence compiled by Conference Board, affects the behaviour of price momentum, another important puzzling in the data that is inconsistent with the CAPM. This pattern has been attributed to market frictions, risk premia and investors' behavioural biases. Providing evidence that clearly support any of these theories resolves the tension around this issue, and helps identify the origins of this puzzling pattern. In this chapter we examine whether price momentum is related to investors sentiment, a behavioural attribute.

The results indicate that momentum is only significant when investors are optimistic, and that these hedge portfolios experience long run reversals. This result is robust to controls to microstructure biases and risk adjustments. It is in line with the behavioural explanation of momentum put forward by Daniel et al (1998), supporting the notion that the foundations of momentum are behavioural.

In the fifth chapter we design an experiment to test the validity of a theory of non-Bayesian updating proposed by Griffin and Tversky (1992). This theory is a cornerstone in the decision making literature as it parsimoniously explains over and under confidence in beliefs, and has been used to explain evidence of over and under reaction in the stock market [Barberis et al (1998)]. A general limitation of studies in psychology, however, is that they do not simulate an economic environment; hence their application in economic models is debatable. Various experiments that have been designed to test the effect of behavioural biases in economic settings have documented mutation or even elimination of the bias (i.e., Grether 1980). Thus, our experiment tests whether this important theory holds for economics decisions, and thus

justifiable used to explain asset prices. Our results suggest that the theory generally holds, but we do document a dramatic decrease of the bias, compared to the original results of GT.

The findings from all empirical chapters demonstrate that systematic deviations from Bayesian updating can occur due to ambiguity aversion, investor sentiment and decision heuristics. Because arbitrage is not limitless [Shleifer and Vishny (1997)] these systematic violations of Bayesian updating can affect asset prices, inducing price predictability. This contradicts the notion that deviations from the Bayesian model are unsystematic, and will thus be washed out during the process of information aggregation [Fama (1998)]. Rather, these results extend the field of behavioural finance by providing new evidence that support the view that price predictability is related to investors' psychology [Hirshleifer (2001)].

The main message that stems from the thesis is that the neoclassical model that assumes rational expectations is incomplete. The behavioural attributes identified in the three empirical chapters of the thesis highlight the incompleteness of the neoclassical model, and suggest that our theories will become more descriptive if we acknowledge the bounded rationality of investors.

6.1 Caveats

As is usually the case in social sciences, the empirical tests conducted in the thesis have limitations. An obvious caveat of the behavioural explanations provided in the third and fourth empirical chapters, is that these patterns in prices may actually reflect compensation for risk bearing, and they only seem anomalous because we have not yet identified the correct equilibrium model of expected returns. This is a common argument proposed by economists who want to dismiss evidence of market inefficiency [see Fama (1998)]. However, in my opinion this

approach, although plausible, is unscientific. Even though it is possible that a rational model may actually explain these patterns, it is not appropriate to dismiss them *a priori* without any strong evidence that such a model exists. Therefore, until this model arrives, results such as these documented in this thesis pose a challenge to the Efficient Market Hypothesis.

In addition, in the fifth chapter the subjects used were university students. Although this is almost the norm in such experiments, there is an issue whether the responses of the students are representative of the entire population, and particularly market participants. If investors learn from their mistakes and gradually become Bayesians the application of such results in financial markets is spurious. However, the evidence on this issue is mixed as various studies have shown that professional investors also exhibit behavioural biases. For example, Coval and Shumway (2005) and Haigh and List (2005) show that futures traders are myopically loss averse. In fact Griffin and Tversky (1992) conclude that experienced subjects are equally likely to exhibit the strength-weight heuristic. Such evidence suggests that it is not conclusive that experienced traders are automatically Bayesians.

Another caveat of the experiment in the fifth chapter is that the parametric methods used to elicit utilities and subjective probabilities are by no means the only ones available. For example, there are various non-parametric alternatives to measure utilities besides the parametric method suggested by Hey and Orme (1994), such as the trade-off method proposed by Wakker and Deneffe's (1996) and the two-step method proposed by Abdellaoui (2000). The main benefit from non-parametric estimations is that the conclusions drawn do not depend on the functional form chosen, which in the fifth chapter it is assumed to be a simple CRRA utility function in the domain of EUT. If, however, subjects make decisions in the risk task according to any other model besides the simple CRRA model our results may lack generality. However, because the Griffin and Tversky (1992) hypothesis involves *comparing* elicited probabilities depending on

the strength and weight of the evidence, the choice of utility model does not seriously limit the generality of our conclusions. This is because, although risk attitude affects elicited beliefs, it is merely an identical transformation of *all* elicited probabilities regardless of the functional form chosen, and thus cannot in principle capture a strength-weight effect. By using a simple parametric CRRA utility function we demonstrate the non-trivial effect of risk attitudes on subjective beliefs, without a significant loss of generality in terms of the conclusions drawn about the Griffin and Tversky (1992) hypothesis.⁷⁵

It is worth briefly expanding on another appealing characteristic of the trade-off method for utility elicitation proposed by Wakker and Deneffe's (1996), apart from being non-parametric. Because it does not involve any deliberation about probabilities it is not sensitive to misspecifications of risk attitudes due to subjects' confusion about probabilities. Our design of eliciting utilities assumes subjects understand and use probabilities correctly, and if they do not produces erroneous estimates of risk attitudes. However, assuming that subjects understand the explicitly given probabilities in the risk task may not be particularly troublesome as they are fairly transparent. For example, as shown in Table 5.1, probabilities correspond to the coloured areas in the pie charts, which are clearly illustrated. In addition, they are well explained by the means of the 100-sided die that we use to decide subjects payments. For example, the probability of receiving £5 in the left panel of Table 5.1 is 87%, which amounts to rolling the 100-sided die and obtaining any number between 1 and 87. This procedure is quite transparent and is unlikely to raise misconceptions that will bias the results. However, having said that, misconceptions are possible, so we raise a flag that our results assume that subjects understand probabilities.

Another caveat of our design is that it does not incorporate probability weighting, as suggested by Kahneman and Tversky (1979) or Quiggin (1992). Although for utility elicitation

⁷⁵ However, the point estimates of the subjective probabilities are sensitive to the model chosen.

the same argument as above still applies, i.e., that a utility model with probability weighting will transform all elicited subjective probabilities in the same manner thus in general will not be able to explain a strength-weight effect, it is possible that probability weighting of *subjective* probabilities is somehow a function of strength and weight. This is an interesting empirical question, which we plan to address in the future.

6.2 Implications for future research

The empirical results in all three chapters imply directions for future research. The third chapter showed that ambiguity aversion affects the pricing of analyst forecasts. However, the literature is lacking a study that examines comprehensively the effect of ambiguity aversion in the cross section of stocks. Perhaps the definition of ambiguity we have provided in this study, namely the richness of the available information, can serve as a basis on which a more general measure of ambiguity can be constructed so that we can examine whether ambiguity affects the cross section of prices.

The fourth chapter demonstrated that momentum is only significant during optimistic periods. Given that the momentum anomaly is correlated with other anomalies, it is interesting to observe what other phenomena investor sentiment can explain. For example, the value premium is negatively correlated with momentum, which suggests that it arises when investors are pessimistic. An investigation of this issue is of great importance as it can highlight whether a puzzling phenomenon such as the value premium is related to sentiment induced mispricing.

The fifth chapter analyzes beliefs from information signals that point to potential gains. However, in financial markets a large fraction of the information set points to contingencies that involve potential losses, and we do not know whether the strength-weight heuristic mutates when

investors are using such signals. Because in the behavioural literature it well documented that “bad is stronger than good” (see Bumeister, Bratslavsky, Finkenauer and Vohs 2001), a mutation of the strength-weight heuristic depending on the domain of the decision is entirely plausible and worth investigating.

In addition, rational decision making does not necessarily imply Bayesian Updating. Rather, in its most general form, it implies that agents make choices according to some “well-behaved” underlying utility function that satisfies the axioms of rational choice. Expected Utility Theory (EUT) is one such decision making model, but assumes that decision makers have complete knowledge of the probabilities that are associated with future contingencies. In our experimental framework subjects *must compute* these probabilities using dice data. Therefore, the framework of their decision is more fitting to a Subjective Expected Utility (SEU) model as outlined by Savage (1972), in which beliefs are idiosyncratic to each individual. Importantly, as explained by Savage, SEU decision makers *need not* form probabilities according to Bayes Rule. All that SEU demands is that the decision-maker behaves *as if* using the mean of the distribution of beliefs, regardless of where these beliefs come from. Savage [1972; p. 57] explains:

According to the personalistic view [of SEU], the role of the mathematical theory of probability is to enable the person using it to detect inconsistencies in his own real or envisaged behaviour. It is also understood that, having detected an inconsistency, he will remove it. An inconsistency is typically removable in many different ways, among which the theory gives no guidance for choosing. Silence on this point does not seem altogether appropriate, so there may be room to improve the theory here.

The psychological heuristic proposed by Griffin and Tversky [1992] can be viewed as one such suggestion to improve the theory, with many other alternatives possible. An interesting question

that merits investigation is to contrast different SEU specifications in terms of the processes that determine the subjective probabilities.

6.3 Closing note

This thesis, as many other studies in this field, demonstrates that the behaviour of investors can actually affect asset prices. This highlights that the neoclassical model, although mathematically very elegant, in reality is very simplistic because it disregards the fact that decisions are a complicated mix of emotions, perceptions and experiences. I would like to close the thesis using a very famous quote from John Maynard Keynes in his 1936 book "*The General Theory of Employment Interest and Money*", which nicely summarizes the "behavioural" nature of decision-making. He notes that:

"Even apart from the instability due to speculation, there is the instability due to the characteristic of human nature that a large proportion of our positive activities depend on spontaneous optimism rather than mathematical expectations, whether moral or hedonistic or economic. Most, probably, of our decisions to do something positive, the full consequences of which will be drawn out over many days to come, can only be taken as the result of animal spirits - a spontaneous urge to action rather than inaction, and not as the outcome of a weighted average of quantitative benefits multiplied by quantitative probabilities." (Keynes 1936, pp 161-162.)

7. Appendices

The appendices aim to expand on certain debates in the field of behavioural finance, so as to give the interested reader a better understanding of the issues in the field. They also provide a literature review on several topics in behavioural finance that are not directly relevant to the main body of the thesis.

7.1 Heterogeneous expectations & behavioural asset pricing

Most asset pricing theories assume that the market has arrived at a *homogenous* Bayesian expectation of the assets risk and return profile, thus determining the fair value of the asset.

Rationality in the marketplace is based on two microeconomic models that have rational (i.e., in the context of the EMH) pricing implications. Firstly, in strict rational expectations suggested by Loucas (1972), investors have perfect foresight; they always interpret information correctly and identically which leads to a Bayesian evolution of prices. This is equivalent to expected utility maximisation as set out by Von-Neumann and Morgenstern (1944). Secondly, rational decision making can occur even if agents do not know the objective probabilities of future contingencies but behave according to the axioms laid out by Savage (1954) in Subjective Utility Theory (SUT). Contrary to EUT, SUT does not require that all investors arrive at identical distributions. They are allowed to have subjective priors which they blend with the new evidence. If this combination occurs according to the rule of Bayes' and under certain conditions (see Blume and Easley 2002),⁷⁶ Subjective Utility Theory (SUT) leads to rational pricing.

This type of rational market is depicted in the diagram below:

⁷⁶ These conditions are that a) the a priori belief includes the true model that generates return or rational priors where the investor has knowledge of the true model and b) the investor correctly anticipates his future beliefs.




Diagram 7.1. Asset market with homogenous expectations.

Suppose that p^* represents the true equilibrium price of an asset at a particular point in time given its risk and return profile.

Formally:

$$p_t^* = d_t E \sum_{t+1}^{t+i} (CF_{t+j}) / \Omega_t \quad (7.1)$$

The price of the asset at time t is the discounted expectation of the future cash flows, CF , given the information available at time t , Ω , where the expectations are rational in the sense of Lucas or Savage. The discount factor depends on the risk properties of the asset's returns. A riskier (less risky) asset will have a larger (smaller) discount factor and thus a lower (higher) price. DD' is the infinitely elastic demand curve for the asset.⁷⁷ This reflects the fact that the market has analyzed the information matrix Ω , and understood its implications on the risk/return profile of the security. Since only one true equilibrium relationship between risk/return can exist given a particular set of fundamentals (properly inferred from information), a horizontal demand curve implies that the market has a homogenous expectation of the cash flows and discount

⁷⁷ Infinitely elastic means that demand at any other price other than p^* demand is 0.

factor, therefore understands that the ‘fair’ price is p^* . Asset demand, therefore, at any other price level above p^* is zero.⁷⁸

This kind of analysis implicitly states that the forces of demand and supply do not affect the asset pricing process, as prices only change in response to changes in fundamentals. For example in diagram 1 if new information in terms of the future cash flows and/or their riskiness arrives in the market, the demand curve would shift to the new implied price level. Individual investors would take positions according to their risk aversion and markets would be at equilibrium. Indeed Ross (1987) makes the point that the forces of demand and supply have little application in financial markets, acknowledging that in neoclassical finance investors *do not speculate*; rather they have complete and correct knowledge of risk/return profiles and simply take positions according to their risk tolerance.

However, the evidence that have accumulated indicate that the functioning of financial markets may not be as in figure 1 and particularly point to investors having heterogeneous beliefs about the ‘fair’ market price. These heterogeneous beliefs may arise because different groups of investors use different information sets to form expectations, as in Hong and Stein (1999) and Peng and Xiong (2006), or because some groups do not process the information as efficiently as other groups (Blume and Easley 2002, Kahneman and Tversky 1982).⁷⁹ Indeed both cases make sense. With the vast amount of assets and information that exists, the complicated interrelationships that pertain within the information set and the fact that attention is scarce, (see Kahneman 1973) it seems unlikely that investors have at their disposal the complete information set when valuing securities. Moreover, the quality of the valuation made by investors depends on the knowledge and experience they possess in the marketplace. Less

⁷⁸ If price is less than p^* the entire market will spot a profit opportunity, rush to buy the stock, thus bid the price up to p^* .

⁷⁹ Brav and Heaton (2002) and Brandt (2003) show that theories based on cognitive biases and theories based on incomplete information are difficult to distinguish empirically.

experienced participants may not be in a position to process the information correctly and may in fact rely on simple heuristics and rules of thumb in order to arrive at decisions.

This evidence suggests that the assumption that markets homogeneously analyze the information set may be flawed as it is conceivable that different groups of investors have different beliefs in terms of what is the 'fair' price of the asset. Furthermore, since at any point given the available information and fundamentals only one risk and return profile describes the company, heterogeneous expectations imply that at least some investors have erred in their valuations.

The issue is then whether these irrational investors affect asset prices. The fact that arbitrage cannot immediately automatically flatten out the demand curve has been discussed in section (2.4). A second aspect to the arbitrage story is that irrational traders will lose money, therefore slowly extinct and prices will be correctly set by the rational investors. This prediction is confirmed by Sandroni (2000) who shows that in a Lucas tree economy only consumers with rational expectations survive. Moreover he shows that if not rational consumers exist, consumers whose forecasts are persistently wrong in the presence of learners, are driven out of the market. Similar results are reached by Blume and Easley (2000) who show that in complete markets where consumers have a common discount factor, those with correct beliefs drive those with incorrect beliefs out of the market and thus steer prices towards fundamentals. The intuition of these models is that agents, being expected (or subjective) utility maximisers allocate their wealth on the events they deem as more likely. Therefore when the true state is revealed only the agents that predicted correctly will accumulate wealth.

This view, however, is challenged on two accounts. Firstly, experimental psychologists challenge the notion that completely rational investors exist as numerous experimental studies find that behavioural biases pertain even when experienced professionals are used (Oskamp

(1965), Baumann, Deber and Thompson (1991), Kidd (1970), Schumway and Coval (2005)). Secondly, a literature has evolved that demonstrates that irrational traders need not be driven out of the market. For example, De Long et al (1990) propose a model with rational arbitragers and irrational traders, and conclude that in equilibrium both types of traders affect returns. Some studies even suggest that irrational traders (through their excessive risk taking) may drive smart money out of the market, (Kyle and Wang (1997), Hirshleifer and Luo (2001), Benos (1998) and Blume and Easley (2002)). Therefore based on this evidence, Friedman's (1953) argument that "dumb" money will eventually be driven out of the market is not uncontroversial.

This implies that demand curves are downward sloping, as in figure 2 below. Prices then become a wealth-weighted average of these expectations and conceivably diverge from fundamentals. For example, suppose that we have two groups of investors, 1 and 2. Due to the reasons highlighted above group 1 believes that the asset should be priced at P1 and group 2 at P2, therefore the demand curve is downward sloping, as shown below:




Diagram 7.2 Market with heterogeneous expectations

The resulting price, given heterogeneous expectations, will be a wealth-weighted average of the valuations of the two groups of investors. If group 1 invests W_1 in the stock and group 2 W_2 then resulting price is given by equation 7.2, as shown below:

$$\hat{P} = \frac{W_1}{W_1 + W_2} P_1 + \frac{W_2}{W_1 + W_2} P_2 \quad (7.2)$$

As this equation shows, the extent that each group affects the stock depends on the relationship between W_1 and W_2 . If the group with the erroneous beliefs, say 2, invests heavily in the stock, its influence, and thus the mispricing will be greater.

Empirical evidence for downward sloping demand curves come from Shleifer (1986). He shows that the stocks that were recently listed in an index, such as the S & P 500, experience a positive price run-up. He says that there exists no reason to expect that the inclusion per se causes a change in fundamentals that warrants this price change. Rather, demand for stocks, due to the various tracker funds, rises after the inclusion of the stock in the index. Shleifer interprets this as evidence of an initially downward sloping demand curve that has shifted outwards.

Similarly, Wugler and Zhuravskaya (2002) state that perfect riskless arbitrage requires the existence of perfect substitute assets. For example the arbitrageur in order to not assume any risk from his arbitrage activities takes two actions. Firstly he trades in a manner that exploits the perceived mispricing. Secondly he performs an equal and opposite act on a perfect substitute so that his position in the market is perfectly hedged. Wugler and Zhuravskaya (2002) state that downward sloping demand curves are a consequence of ‘arbitrage risk’ due to the lack of perfect substitutes. In effect this paper confirms the theoretical proposition of Shleifer and Vishny (1997).

These findings are in direct contradiction with Ross's claim that demand and supply do not affect market prices. If a group of investors, for example, increases its demand for a security for whatever reason and if the amount invested is large enough the price of the asset will change. Therefore prices not only change as a response to changes in fundamentals but rather they change because a particular group has changed its valuation of the asset. This opens the door to a completely new perspective that allows prices not only to equal rationally discounted cash flows but to reflect *expectational equilibria*, [Kandel and Pearson (1995), Barberis, Shleifer and Vishny (1998), La Porta (1996)], thus justifying "behavioural" asset pricing.

7.2 Do investors systematically over or under react?

On the theoretical level researchers have identified various factors that cause investors to systematically over or under react [e.g. Barberis et al. (1998), Daniel et al (1998) and Odean (1998)]. If these models are correct and return reversals and continuations are due to *systematic* misperceptions of information, empirical studies should be able to demonstrate that, in response to particular information signals, prices either trend or revert. This is important because ultimately empirical work will determine whether the behavioural explanation of return predictability is valid. As noted by Barberis and Thaler (2002 pp.61) “There is only one scientific way to compare alternative theories, behavioural or rational, and that is with empirical tests.”

Fama (1998) challenges the validity of behavioural models as he reviews the literature on return predictability from past information and concludes that *overreaction is as likely as underreaction*; he argues therefore that behaviour is randomly split between these models of behaviour, thus behavioural finance does not yield any valuable insights. In order for behaviouralists to overcome Fama’s critique they must identify parsimoniously the conditions that spur over and under reactions amongst investors, as predicted by the theoretical models.

We identify three studies in the literature that explicitly attempt to examine when investors may over or under react. Chan, Kothari, and Frankel (2004) review the behavioural literature and suggest that the phenomena of conservatism (Edwards 1968) and representativeness (Kahneman 1974) are key ingredients of behavioural explanations. Their attempt is to construct measures of past company performance, which will serve as the information signals that investors misperceive. Particularly they state that investors will overreact, due to representativeness, to companies which have experienced consistent positive or

negative performance. On the contrary, due to conservatism, when a firm has a short period of good or bad performance, investors will underreact and cause returns to trend. The authors use 5 and 1 year measures based on sales and income per share. They conclude that no evidence of return reversals or return continuations is found.⁸⁰ The authors state that this evidence does not support the behavioural model. They state however, that these results apply only to their sample and their calculated measures; therefore they do not preclude the existence of these biases in the marketplace under *different* conditions.

Kadiyala and Rau (2004) examine whether the over or under reaction models of behaviour fits the corporate events of seasoned equity offerings, stock/cash financed acquisitions and share repurchases. They suggest that these events can be classified to good or bad information. For example share repurchases and cash finance acquisitions convey good news and SEO's and stock financed acquisitions bad news. They also control for the earnings surprise prior to such events in an attempt to see whether people overreact when bad performance (i.e. negative surprise) is coupled with bad news and vice versa. Similarly to Chan et al (2004) they only find return drifts which depend mainly on the earnings surprise. Therefore they do not identify conditions that may drive investors to under or over react.

An innovative paper that provides evidence of both return reversals and continuations is Chan (2003). He uses a sample in which he distinguishes stocks 'with news' (any news in the media) and 'no news'. He then scrutinizes the data to many different tests and concludes that the 'with news' stocks experience drifts and the 'no news' stocks, which have experienced large price swings, undergone return reversals. This paper is a major breakthrough because the author produces these opposite return patterns in the same data set, which suggests that these behaviors may not be random, as suggested by Fama (1998). However it makes no predictions of the

⁸⁰ The authors state that some evidence for return continuations are found, which however are sensitive to the post-earnings announcement drift.

conditions that drive over or under reaction. For example the stocks that experience large price swings and subsequent reversals probably involve some kind of news, which the study did not or could not observe. Therefore, an open question remains in terms of identifying factors that may drive investors to systematically over or under react.

7.3 Further literature on market anomalies

The purpose of the appendix is to provide more details for some important market anomalies that have been associated with investors' behavioural biases.

7.3.1 Price Momentum

As explained, Jegadeesh and Titman (1993) document the phenomenon of price momentum. Stocks that performed well in the recent past (winners) continue to outperform stocks that performed poorly. Jegadeesh and Titman (1993) show that a 'momentum' strategy that is long in past winners and short in past losers earns a substantial abnormal return. The robustness of this strategy has been established in different markets and is currently being widely exploited by the investment community.

Possible explanations for this anomaly are a) slow information diffusion, b) 'underreaction' and c) positive feedback trading. In terms of the former Hong and Stein (1999) construct a model where information diffuses gradually in the market. Some groups may have access to the same information prior to other groups. This creates non-synchronous trading effects that show up as momentum in prices. Hong and Stein (2000) empirically confirm this prediction. Therefore it seems logical that to an extent the phenomenon of momentum depends on market frictions.

The behavioural explanation of the phenomenon of momentum is that investors systematically exhibit conservatism-driven behaviour when assessing news. This implies that since, information is 'irrationally' discounted the time period required for its implications to be reflected in asset prices is prolonged. This translates to a drift, as opposed to a jump, towards fundamentals. In a very important paper Zhang (2006) shows that the momentum increases in information uncertainty. He argues that if investors systematically display conservatism in

general they will do so even more when the information is of lower quality. He forms two- way portfolios based on momentum and various measures of information uncertainty, such as size, age, analyst coverage, volatility, and shows that the returns increase in information uncertainty (IU), and shows that momentum increases with IU. These results support the underreaction story.

In terms of positive feedback trading Chan, Jegadeesh and Lakonishok (1996) rigorously test a number of hypotheses that aim to determine whether momentum is a consequence of an underreaction to public information, such as earnings and revisions in analyst forecasts, (earnings momentum) or to a series of private information (price momentum), or to positive feedback trading. They find no evidence for the latter as no reversals actually occur. In terms of earnings and price momentum the authors find that no strategy subsumes the other as they both capture different effects. The authors conclude that investors underreact to both public information for short- term earnings as well as private information that relates to long-term profitability.

7.3.2 The B/M effect

An anomaly observed in the marketplace is that the Book to Market ratio explains average returns. Lakonishok, Shleifer and Vishny (1994) state that high B/M (value) stocks systematically outperforms low B/M (growth) because investors suffer from the extrapolation bias (representativeness heuristic). Growth stocks entail good news, which however are falsely extrapolated in the future. Investors become too optimistic for these stocks leading to over pricings. The converse holds for value stocks. The fact that value predictably outperforms growth, according to Lakonishok et al (1994), reflects the correction process towards the true fundamental values. Supportive evidence for this come from La Porta, Lakonishok, Shleifer and Vishny (1997) as they show that the surprise after earnings announcements is systematically more positive for value, as opposed to growth stocks, reflecting that investors were pessimistic

for the former and optimistic for the latter. Daniel and Titman (2006) decompose the B/M factor to a backward and forward looking component and show that the optimism, which generates the phenomenon, stems from an overreaction to the forward looking component. However Fama and French (1993) argue that the B/M factor captures financial distress and is thus a source of undiversifiable risk. Fama and French (1996) extend the CAPM to an ICAPM with three factors, a size (SMB) and a B/M (HML) based factors. Petcova (2006) offers support that the HML and the SMB factors correlate with proxies for changes in investment opportunities, thus justifying the ICAPM explanation given by Fama and French (1996) to their three-factor model. Petcova and Zhang (2005) show that value is riskier than growth, therefore claim that the out performance of the former does not constitute irrationality. In addition Doukas, Kim and Pantzalis (2002) test the error in expectations hypothesis proposed by Lakonishok et al (1994) and find contradictory evidence.

7.3.3 Post-Earnings announcement drift

The seminal paper of Ball and Brown (1978) demonstrated a drift in security prices after earnings announcements. Returns for companies that have announced earnings above (below) market expectations exhibit a positive (negative) momentum. These results have spawned the appearance of a massive literature. A representative study, for example Foster, Olsen and Shevlin (1984), finds the drift pertains 60 days after the announcement and a strategy that's aims to exploit it, by taking a long position in the 'good news' portfolio and short in the 'bad news' portfolio, makes a substantial abnormal annual return of 25%. Similar results are reported by Bernard and Thomas (1989, 1990), Ball and Kothari (1991), Rendleman, Jones and Latane

(1982) and more recently Chordia and Shivakumar (2005), Doyle, Lundholm and Soliman (2003), Livnat (2003), Garfinkel and Sokobin (2005).

The robustness of the profitability of this trading strategy across markets and time periods is a major source of debate amongst academics. Eugene Fama, the most influential supporter of Efficient markets, has referred to the post earnings announcement drift (PEAD) as the ‘father of all anomalies’.

Why are these findings troublesome? In the light of the efficient market hypothesis, as put forward by Fama (1970), any kind of predictability is impossible. Prices are quick to respond *correctly* to information, thus maintain the martingale property. However the PEAD literature posits a clear violation in that prices are predictable based on whether the earnings of a company are below or above market expectations.

The model that is universally used to sort stocks in ‘good’ and ‘bad’ news portfolios in all studies that document the PEAD is the Standard Unexpected Earnings model (SUE),

$$\frac{E_t^A - E_t^e}{\text{Deflator}}$$
 where E^A corresponds to the actual earnings per share reported at time t, E^e is a measure of what the market expects earnings to be at time t and the denominator is a standardising deflator, for example stock price at time t. The literature documents a monotonic relationship between abnormal returns and SUE. The higher (lower) SUE decile experiences the highest (lowest) cumulative abnormal return in the window after the announcement.

In terms of the PEAD literature a considerable effort is devoted in developing a more precise measure for the market’s expectations. Finding such a measure will enable the construction of more informative ‘surprise’ portfolios and hence a better understanding, and exploitation, of the PEAD. In the early stages of this literature research utilized findings from the accounting literature, Brown (1993), and used earnings forecasts from time series models a

proxy for the market's expectations (Foster (1977), Bartov (1992), Griffin (1977) Bernard and Thomas (1990). As the literature developed further, researchers use the consensus analyst forecast of the EPS as the markets expectation for the upcoming EPS (Fried and Givoly 1982, Givoly and Lakonishok (1979), Elton, Gruber and Gultekin (1981), Brown, Foster and Noreen (1984), Livnat and Mendenhall (2006). Using analyst forecasts as expectations provides a better fit to the data. As documented by Livnat and Mendenhall (2006), the PEAD is stronger when expectations are defined using the I/B/E/S analyst forecasts in comparison to a time series model, suggesting that market expectations are better captured by analyst forecasts.

The general consensus in terms of the PEAD is that investors *underreact* when a company reveals higher or lower earnings than expected. Therefore, prices slowly drift toward fair values, generating the drift.

7.3.4. Individual trading

The most convincing evidence for the existence of behavioural biases come from a series of papers by Odean. He uses a unique data set that includes the trades of individual investors. He can thus identify which stocks were bought and sold at which point in time, therefore his sample allows inference in terms of whether these investors are rational. Odean (1999) shows that investors trade too much and suffer capital losses in the process. Neoclassical finance assumes that investors have rational expectations. This implies that investors would only trade if the marginal benefit from doing so exceeds the marginal cost, Grossman and Stiglitz (1980). Odean finds that the securities that are bought do not outperform the securities that were sold by enough to cover the costs of trading. Moreover, and much more worryingly for the case of neoclassical finance, the securities that were sold, on average, *outperform* the securities that were bought.

This implies that the information set was completely misinterpreted by investors and that their expectations were largely biased. For example suppose that the return of a stock will be 2% in the following month, implying that the news for this particular company at this moment is good. The rational agent will form an expectation of 2%. However an agent that forms an expectation of 1.7% is biased in a neoclassical sense but in reality he is *quasi-rational* (i.e. does not err substantially). He understands that the information is positive, but misjudges the magnitude. If however the investor believes that the stock will lose -1.5%, as appears to be the case in Odean's sample, the investor was unable to understand the nature of the information set.

Barber and Odean (2000) partition this sample in quintiles based on trading activity and test the common stock performance of the different groups. They find that the group that trades the most makes the lowest return, a finding that again contradicts the argument of Grossman and Stiglitz (1980). Further, Odean and Barber (2002) show that the traders that switch from phone to online accounts trade much more and make lower returns. The finding from this study is clear and puts into perspective the findings of the other two papers. Traders, at least the 'amateurs', simply do not trade on fundamentals. Rather they speculatively trade on noisy information signals that they fail to interpret correctly. Consequently their investment performance suffers. Odean attributes these findings to the overconfidence bias. He states that investors believe that they are better at selecting securities than what they really are, thus trade more frequently not on substantive information but on noise. The investment information set is a very noisy and dynamic platform, and this research demonstrates that agents, being computationally bounded, cannot form unbiased expectations.

Odean (1998) tests the disposition effect, put forward by Shefrin and Statman (1985), which states that investors are quick to sell the winning stocks in order to realize the gains and ride losers to long as to avoid realizing the losses. This type of behaviour stems from Kahneman

and Tversky's (1979) Prospect Theory. Odean (1998) finds that the investors in his sample exhibit this kind of behaviour. Winner stocks are sold very quickly, thus failing to capture future gains, and loser stocks are sold very late, thus exacerbating the loss. This is evidence that individual investors *fail* to form *unbiased* expectations.

Odean and Barber (2005) show that individual traders are prone to 'attention' trading and that they are more likely to buy stocks which they owned previously. As a result, Odean and Barber (2005) show that individual traders buy high and sell low, which again contradicts rational expectations.

The research of Odean is really the most conclusive in terms of investor behaviour being bounded as he clearly documents deviations from neoclassical theories. A question that remains is whether these behavioural biases affect asset prices. The investors in Odean sample invest little capital, about \$11,000 per year therefore despite their biased behaviour, prices may not be affected. Moreover the investors in Odean and Barber's sample, since they invest little capital, are most likely small individual investors. Therefore these findings cannot be expected to apply to large institutional investors. Some evidence exists that support this conjecture that the behaviour of professionals is quasi-rational. In a sample that includes both individual and institutional investors Malmendier and Shanthikumar (2009) show that the behaviour of the former is bias-prone whereas the latter appear to be rational. This suggests large investors, being efficient processors of information, should be able to spot mispricings and arbitrage them away.

However, evidence that behavioural biases pertain amongst experienced people also exist. For example, Coval and Shumway (2005) use a data set that contains the trades of professional future traders amongst the Chicago Board of Trades. The authors state that these traders exchange 200 million dollars worth of future contracts daily. Thus any behavioural biases in this group are likely to affect prices. The authors find that these traders are significantly loss-

averse, as subsequently to losses their risk taking increases significantly so that they recover their position. In terms of loss-aversion affecting prices the authors find that the market is able to identify the informed from the loss-averse trades; hence prices for the latter revert much more rapidly to their fundamental values. This finding highlights the capacity of the market to fill mispricings and supports that ideas of Friedman (1953) and Fama (1980), that biased trades will eventually be driven out of the market. However, these results also show that prices, albeit temporarily, are influenced by behavioural trades so it is conceivable that in an environment where 'behavioural volume' is greater the effects on prices are long lasting. In any case the results of Coval and Shumway (2005) leave an open question as to whether the behavioural biases, documented by Odean, substantially prohibit market efficiency.

References

- Abarbanell S. J., Lanen N. W., Verrecchia E. R, 1995, “Analysts’ Forecasts as Proxies for Investor Beliefs in Empirical Research,” *Journal of Accounting and Economics*, 20, 31-60.
- Abdellaoui M., 2000, “Parameter-free elicitation of utility and probability weighting functions”, *Management Science*, 46, 1497-1512.
- Acemoglu D., Scott A., 1994, “Consumer Confidence and rational expectations: Are agents beliefs consistent with the theory?,” *The Economic Journal*, 104, January, 1-19.
- Ahn D., Choi S., Gale D., Karin S., 2009, “Estimating ambiguity aversion in a portfolio choice experiment,” *UC Berkley Working Paper*.
- Andersen S., Fountain J., Harrison W., G., Rutström E. E., 2007, “Eliciting Beliefs: Theory and Experiments,” *Working Paper 07-08*, Department of Economics, College of Business Administration, University of Central Florida.
- Anderson T. W., 1984, “An introduction to multivariate statistical analysis”, Wiley series in Probability and Mathematical statistics.
- Arkes H., Herren L., Isen A., 1988, “The role of potential loss in the influence of affect on risk taking behavior,” *Organizational Behavior and Human Decision Processes*, 35, 124-140.
- Arrow J. K., 1986, “Rationality of self and others in an economic system,” *Journal of Business*, 59, 385-399
- Asness C., Moskowitz T., Pedersen L., 2009, “Value and Momentum Everywhere,” Working paper, New York University.
- Badrinath S.G., Wahal S., 2002, “Momentum Trading by institutions,” *Journal of Finance*, 57, 2449-2478.
- Baker M., Wurgler J., 2006, “Investor Sentiment and the Cross-Section of Stock Returns,” *Journal of Finance* 61, 1645–1680.
- Baker, M., Wurgler, J., 2007, “Investor Sentiment in the Stock Market,” *Journal of Economic Perspectives* 21, 129-152.
- Ball R., Brown P., , 1968, “An empirical evaluation of accounting income numbers”, *Journal of Accounting Research*, 6, 159–178.
- Ball R., Kothari S. P., 1991, “Security returns around earnings announcements”, *Accounting Review*, 66, 718-738.
- Banz W. R., 1981, “The relationship between return and market value of common stocks,” *Journal of Financial Economics*, 9, 1, 3-18.

- Barber M. B., Loeffler D., 1993, "The Dartboard column: Second-hand information and price pressure," *Journal of Financial and Quantitative Analysis*, 28, 2, 273-284.
- Barber B. M., Lyon J. D., 1997, "Detecting long run abnormal returns: The empirical power and specification of test statistics", *Journal of Financial Economics*, 43, 341-372.
- Barber M. B., Odean T., 2000, "Trading is hazardous to your health: The common stock investment performance of individual investors", *Journal of Finance*, 55, 773-806.
- Barber M. B., Odean T., 2001, "Boys will be boys: Gender, Overconfidence, and common stock investment", *Quarterly Journal of Economics*, 116, 261-292.
- Barberis N., Huang M., Santos T., 2001, "Prospect Theory and asset prices", *Quarterly Journal of Economics*, 116, 1-53.
- Barberis N., Huang M., 2008, "Stocks as lotteries: The implication of probability weighting for security prices", *American Economic Review*, 98, 2066-2100.
- Barberis N., Shleifer A., Vishny R., 1998, "A model of investor sentiment," *Journal of Financial Economics*, 49, 307-343.
- Barberis N., Thaler R., "A survey of behavioural finance", IN Handbook of the economics of Finance, 1, 1053-1128.
- Bartov, E., 1992. Patterns in unexpected earnings as an explanation for post-announcement drift. *The Accounting Review* **67**, pp. 610–622.
- Baumann AO, Deber RB, Thompson GG., 1991, "Overconfidence among physicians and nurses: the 'micro-certainty, macro-uncertainty' phenomenon.", *Social Science and Medicine*, 32(2):167-74.
- Baumeister R. F., Bratslavsky E., Finkenauer C., Vohs K. D., 2001, "Bad is stronger than good", *Review of General Psychology*, 5, 323-370.
- Becker S., W., Brownson F. O., 1964, "What price ambiguity? Or the role of ambiguity in decision making", *Journal of Political Economy*, 72, 62-73.
- Bekaert G., Wu G., 2000, "Asymmetric volatility and risk in equity markets," *Review of Financial Studies*, 13, 1-42.
- Benos, A. V., 1998, "Aggressiveness and survival of overconfident traders", *Journal of Financial Markets*, 1 (3), 353-383.
- Bergman N., Roychowdhury S., 2008, "Investor sentiment and corporate disclosure," *MIT working paper*.

Bernard V. L., Thomas J. K., 1989, "Post-Earnings announcement drift: Delayed price response or risk premium?", *Journal of Accounting Research*, 27, 1-36.

Bernard V., Thomas J. K., 1990, "Evidence that stock prices do not fully reflect the implications of current earnings for future earnings", *Journal of Accounting and Economics*, 13, 305-340.

Bhushan R., 1989, "Firm characteristics and analyst following", *Journal of Accounting and Economics*, 11, 255-274

Bier V. M., Connell B. L., 2006, "Ambiguity seeking in multi-attribute decisions: Effects of optimism and message framing", *Journal of Behavioral Decision Making*, 7, 169-182.

Binmore K., 2008, "Rational Decisions," Princeton University Press.

Black F., 1986, "Noise," *Journal of Finance*, 41, 529-543.

Blume, L. E. & Easley, D., 2000, "If you are smart why you aren't rich? Belief selection in complete and incomplete markets", unpublished manuscript, Cornell University.

Blume, Lawrence E. & Easley, David, 2002. "Optimality and Natural Selection in Markets," *Journal of Economic Theory*, vol. 107(1), pages 95-135.

Bonner E. S., Walther R. B., Young S. M., 2003, "Sophistication-Related differences in investors models of the relative accuracy of analysts forecast revisions," *The Accounting Review*, 78, 3, 679-706.

Bower G., 1981, "Mood and Memory," *American Psychologist*, 36, 129-148.

Bower G., 1991, "Mood congruity of social judgment." In J. Forgas (Ed.) *Emotion and Social Judgment*, Oxford, England: Pergamon Press.

Bracha, A., Brown, J. D., 2009, "Affective Decision Making: A Behavioural Theory of Choice", *Cowles Foundation working paper*, Yale University.

Brav A., Heaton J. B., 2002, "Competing theories of financial anomalies", *Review of Financial Studies*, 15, 575-606.

Brown, L., 1993. Earnings forecasting research: its implications for capital markets research. *International Journal of Forecasting* 9, pp. 295–320.

Brown L., 2001, "How Important is Past Analyst Forecast Accuracy?," *Financial Analysts Journal*, November, 44-49.

Brown W. G., Cliff, M. T., 2005, "Investor sentiment and asset valuation," *Journal of Business*, 78, 405-439.

Brown, P., Foster, G. and Noreen, E., 1985. *Security Analyst Multi-Year Earnings Forecasts and the Capital Market*, American Accounting Association, Sarasota, FL.

- Brown D. L., Rozeff, M. S., 1980, "Analysts can forecast accurately!," *Journal of Portfolio Management*, 6, 31-34.
- Brown S. J., Warner J. B., 1985, "Using daily stock returns", *Journal of Financial Economics*, 14, 3-31.
- Brunnermeier M. K., Nagel S., 2005, "Hedge funds and the technology bubble", *Journal of Finance*, 59, 2013-2040.
- Butler, K. C., Lang, L.H., 1991, "Thye forecast accuracy of individual analysts: Evidence of systematic optimism and pessimism," *Journal of Accounting Research*, 29, 150-156.
- Camerer C., Weber M., 1992, "Recent Developments in Modelling Preferences: Uncertainty and Ambiguity," *Journal of Risk and Uncertainty*, 5, 325-370.
- Camerer F., C., Kunreuther H., 1989, "Experimental Markets for Insurance," *Journal of Risk and Uncertainty*, 2, 265-300.
- Campbell J. Y., Hilscher J., Szilagyi J., "In search of distress risk", *Journal of Finance*, forthcoming.
- Campbell J. Y., Lo A. W., MacKinlay A C., 1997. *The Econometrics of financial markets*. Princeton University Press. Princeton. New Jersey.
- Capstaff J., Paudyal K., Rees W., 1995, "The Accuracy and Rationality of Earnings Forecasts by UK Analysts, *Journal of Business, Finance and Accounting*, 22, 67-85.
- Cao H. H., Wang T., Zhang H. H., 2005, "Model Uncertainty, limited participation, and asset prices", *Review of Financial Studies*, 2005, 18, 1219-1251.
- Chan W.S., 2003, "Stock price reaction to news and no-news: drift and reversal after headlines", *Journal of Financial Economics*, 70, 223-260.
- Chan, W., Kothari, S., Frankel, R., 2004, "Testing Behavioural Finance Theories Using Trends and Consistency in Financial Performance", *Journal of Accounting and Economics* 38, 3-50.
- Chan K. L., Jegadeesh N., Lakonishok J., 1996, "Momentum Strategies," *Journal of Finance*, 51, 1681-1713.
- Chen Z., Epstein L., 2002, "Ambiguity, Risk and Asset Returns in Continuous Time", *Econometrica*, 70, 4, 1403-1443.
- Chordia T., Shivakumar L., 2002, "Momentum, Business cycle, and time varying expected returns," *Journal of Finance*, 57, 985-1019.

- Chordia T., Shivakumar L., “Earnings and price momentum”, *Journal of Financial Economics* 80, 627–656.
- Chordia T., Roll R., Subrahmanyam A., 2001, “Market liquidity and trading activity,” *Journal of finance*, 56, 501-530.
- Chordia T., Roll R., Subrahmanyam A., 2002, “Order imbalance, liquidity, and market returns,” *Journal of Financial Economics*, 65, 111-130.
- Clement M., 1999, “Analyst forecast accuracy: Do Ability, Resources and Portfolio Complexity Matter?”, *Journal of Accounting and Economics*, 27, 285-303.
- Clement M., Tse S., 2003, “Do Investors Respond to Analysts’ Forecast Revisions as if Accuracy is All that Matters?,” *The Accounting Review*, 78, 227-249.
- Cohen M., Gilboa I., Jaffray Y., J., Schmeidler D., 2000, “An Experimental Study of Updating Ambiguous Beliefs,” *Risk Decision and Policy*, 5, 123-133.
- Conrad J., Cornell B., Landsman W., 2002, “When is Bad News Really Bad News?,” *Journal of Finance*, 57, 2507-2532.
- Conrad J., Kaul G., 1998, “An anatomy of trading strategies,” *Review of Financial Studies*, 11, 489-519.
- Cooper M. J., Gutierrez C. Jr., Hameed A., 2004, “Market states and momentum,” *Journal of Finance*, 49, 1345-1365.
- Cooper M. J., Dimitrov O., Rau R. P., 2001, “A rose.com by any other name,” *Journal of Finance*, 56, 2371-2388.
- Cornelli F., Goldreich D., Ljungqvist A., 2006, “Investor sentiment and Pre-IPO markets,” *Journal of Finance*, 61, 1187-1216.
- Coval J., Shumway T., 2005, “Do Behavioral Biases Affect Prices?,” *Journal of Finance*, 60, 1,1 – 34.
- D’Avolio G., 2002, “The market for borrowing stock,” *Journal of Financial Economics*, 66, 271-306.
- Daniel K., Titman S., 1997, “Evidence on the characteristics of cross sectional variation in stock returns,” *Journal of Finance*, 52, 1-33.
- Daniel K., Grinblatt M., Titman S., Wermers R., 1997, “Measuring mutual fund performance with characteristic-based benchmarks,” *Journal of Finance*, 52, 1035-1058.
- Daniel K., Hirshleifer D., Subrahmanyam A., 1998, “Investor psychology and security market under-and overreactions,” *Journal of Finance*, 53, 1839-1885.

- Daniel K., Titman S., 2006, "Market reactions to tangible and intangible information", *Journal of Finance*, 61, 1605-1643.
- De Bondt W. F. M., Thaler R., 1985 "Does the stock market overreact?," *Journal of Finance*, 40, 793-805.
- De Bondt W. F. M., Thaler R., 1987, "Further evidence on investor overreaction and stock market seasonality?," *Journal of Finance*, 42, 557-581.
- Dechow P., Dichev I.D., 2002, "The quality of earnings: The role of accrual estimation errors", *The Accounting Review*, 77, 35-59.
- De Long, A. Shleifer, L. Summers and R. Waldman, 1990, "Noise trader risk in financial markets", *Journal of Political Economy*, 98, 703-738.
- Diether K., B., Malloy C. J., Scherbina A., 2002, "Differences of opinion and the cross section of stock returns", *Journal of Finance*, 57, 2113-2141.
- Doukas J. A., Kim C., Pantzalis C., 2002, "A test of the errors in expectations hypothesis: Evidence from analyst forecasts", *Journal of Finance*, 57, 2143-2165.
- Doukas A. J., McKnight J. P., 2005, "European momentum strategies, information diffusion and investor conservatism," *European Financial Management*, 11, 313-338.
- Doyle J., R. Lundholm and M. Soliman, "The predictive value of expenses excluded from pro forma earnings", *Review of Accounting Studies* 8, 145-174
- Dreman D., 1998, *Contrarian investment strategies: the next generation: beat the market by going against the crowd.* Simon and Schuster. New York.
- Edmans A., Garcia D., Norli Ø., 2007, "Sports sentiment and stock returns," *Journal of Finance*, 62, 1967-1998.
- Edwards W., "Conservatism in human information processing", in B. Kleinmuntz (Ed.), *Formal representation of human judgment.* New York: Wiley, 17-52.
- Einhorn J. H., Hogarth M. R., 1985, "Ambiguity and Uncertainty in Probabilistic Inference," *Psychological Review*, 92, 4, 433-462.
- Ellsberg D., 1961, "Risk, Ambiguity and the Savage Axioms", *Quarterly Journal of Economics*," 75, 643-669.

Elton, E., Gruber, M. and Gultekin, M., 1984. Professional expectations: accuracy and diagnosis of errors. *Journal of Financial and Quantitative Analysis* 19, pp. 351–363

Epstein L., Schneider M., 2003, “Recursive multiple priors,” *Journal of Economic Theory*, 113, 1-31.

Epstein L., Schneider M., 2008, “Ambiguity, Information Quality and Asset Pricing,” *Journal of Finance*, 63, 1, 197-228.

Fama E. F., 1979, “Efficient capital markets: A review of theory and empirical work”, *Journal of Finance*, 383-417.

Fama E., 1985, “The behaviour of stock market prices,” *Journal of Business*, 38, 34-105.

Fama E., 1998, “Market efficiency, long run returns and behavioural finance”, *Journal of Financial Economics*, 49, 283-306

Fama E., French K., 1988, “Permanent and temporary components of stock prices,” *Journal of Political Economy*, 96, 246-273.

Fama E., French K., 1992, “The cross section of expected stock returns,” *Journal of Finance*, 47, 427-465.

Fama E., French K., 1993, “Common risk factors in the returns of stocks and bonds,” *Journal of Financial Economics*, 33, 3-56.

Fama E., French K., 1995, “Size and Book to Market Factors in Earnings and Returns,” *Journal of Finance*, 50, 1, 131-155.

Fama E., French K., 1996, “Multifactor explanations of asset pricing anomalies,” *Journal of Finance*, 51, 55-84.

Fehle F., Tsyplakov S., Zdorovtsov M. V., 2005, “Can Companies Influence Investor Behaviour through Advertising? Super Bowl Commercials and Stock Returns”, *European Financial Management*, 11, 5, 625-647.

Finucane L. M., Alhakami A., Slovic P., Johnson M. S., 2000, “The affect heuristic in judgments of risks and benefits,” *Journal of Behavioural Decision Making*, 13, 1-17.

Fiore M. S., Harrison W. G., Hughes E. C., Rutström E. E., 2008, “Virtual Experiments and Environmental Policy”, *Journal of Environmental Economics and Management*, forthcoming.

Fishburn P. C., 1988, *Nonlinear preference and utility theory*. Baltimore: John Hopkins University Press.

- Fisher L. K., Statman M., 2003, "Consumer confidence and stock returns," *The Journal of Portfolio Management*, 30, 115-127.
- Fitzsimons G., Shiv B., 2001, "Non-conscious and Contaminative Effects of Hypothetical Questions on Subsequent Decision Making," *Journal of Consumer Research*, 28, 224-238.
- Fiore M S., Harrison W. G., Hughes E. C., Rustron E. E., 2009, "Virtual experiments and environmental policy," *Journal of Environmental Economics and Management*, 57, 65-86.
- Foster, G., 1977. Quarterly accounting data: time-series properties and predictive-ability results. *Accounting Review* 52, pp. 1–21.
- Foster G., Olsen C., Shevlin T., 1984, "Earnings releases, anomalies and the behaviour of security returns", *Accounting Review*, 59, 574-603.
- Frazzini A., "The disposition effect and underreaction to news", *Journal of Finance*, 61, 2017-2046.
- Frazzini A., Lamont A. O., 2008, "Dumb Money: Mutual fund flows and the cross section of stock returns," *Journal of Financial Economics*, 88, 299-322.
- Fried, D. and Givoly, D., 1982. Financial analysts' forecasts of earnings: a better surrogate for market expectations. *Journal of Accounting and Economics* 4, pp. 85–107
- Frisch D., and Baron J., 1988, "Ambiguity and Rationality," *Journal of Behavioral Decision Making*, 1, 149 -157.
- Friedman M., 1953, "The methodology of positive economics," in *Essays in Positive Economics*, Chicago University Press
- Fuller K., Netter J., Stegemoller M., 2002, "What do returns to acquiring firms tell us? Evidence from firms that make many acquisitions", *Journal of Finance*, 57, 1763-1793.
- Gärdenfors P., Sahlin N. E., 1983, "Decision Making with Unreliable Probabilities," *British Journal of Mathematical and Statistical Psychology*, 36, 240-251.
- Gärderfors P., Sahlin N. E., 1982, "Unreliable Probabilities, risk taking, and decision making," *Synthese*, 53, 361-386.
- Garfinkel, J., & Sokobin, J. (2006). Volume, opinion divergence, and returns: A study of post-earnings announcement drift. *Journal of Accounting Research*, 44, 85–112.
- Gilboa I., Schmeidler D., 1989, "Maxmin Expected Utility With Non-Unique Prior," *Journal of Mathematical Economics*", 18, 141- 153.

- Givoly D., Lakonishok J., 1979, "The Information Content of Financial Analysts' Forecasts of Earnings," *Journal of Accounting and Economics*, 1, 165-185.
- Glaser M., Weber M., 2009, "Which past returns affect trading volume?", *Journal of Financial Markets*, 12,1-31.
- Gleason C., Lee C., 2003, "Analyst Forecast Revisions and Market Price Discovery", *The Accounting Review*, 78, 1, 193-225.
- Goetzmann N. W., Massa M., 2002, "Daily momentum and contrarian behavior of index fund investors," *Journal of Financial and Quantitative Analysis*, 37, 325-389.
- Grether D., 1980, "Bayes Rule as a Descriptive Model: The Representativeness Heuristic," *Quarterly Journal of Economics*, Vol.95, No. 3, 537-557.
- Grether D., Plott R. C., 1979, "Economic Theory of Choice and the Preference Reversal Phenomenon," *American Economic Review*, September 1979, 623-638.
- Griffin, P., 1977. The time-series behavior of quarterly earnings: preliminary evidence. *Journal of Accounting Research* **15**, pp. 71–83
- Griffin D., Tversky A., 1992, "The Weighing of Evidence and the Determinants of Confidence," *Cognitive Psychology*, 24, 411-435.
- Griffin M. J., Ji X., Martin S. J., 2003, "Momentum Investing and business cycle risk: Evidence from pole to pole," *Journal of Finance*, 58, 2515-2547.
- Grinblatt M., Han B., 2005, "Prospect theory, mental accounting and momentum," *Journal of Financial Economics*, 78, 311-339.
- Grinblatt M., Keloharju M., 2000, "The investment behavior and performance of various investor types: A study of Finland's unique data set," *Journal of Financial Economics*, 55, 43-67.
- Grinblatt M., Keloharju M., 2009, "Sensation seeking, overconfidence and trading activity" *Journal of Finance*, 64, 549-578.
- Grossman S. J., Stiglitz J. E., 1980, "On the impossibility of informationally efficient markets", *American Economic Review*, 70, 393-408.
- Grundy B. D., Martin J. S., 2001, "Understanding the nature of the risks and the source of the rewards to momentum investing," *Review of Financial Studies*, 14, 29-78.
- Haigh M. S., List J. A., 2005, "Do professional traders exhibit myopic loss aversion? An experimental analysis", *Journal of Finance*, 60, 523-534.

- Harrison W. G., 2006, "Hypothetical Bias over Uncertain Outcomes," in J. List (Ed.), *Using experimental methods in environmental and resource economics* (Northampton, MA: Elgar).
- Harrison W. G., 1994, "Expected Utility Theory and the Experimentalists," *Empirical Economics*, 19, 223-253.
- Harrison W. G., Johnson E., McInnes M. Melayne, Rutström E. Elisabet, 2005, "Risk Aversion and Incentive Effects: Comment," *American Economic Review*, 95, 3, 897-901.
- Harrison W. G., Rutström E. E., 2008, "Risk Aversion in the Laboratory," in J.C. Cox and G.W. Harrison (Eds.), *Risk Aversion in Experiments* (Bingley, U.K: Emerald, Research in Experimental Economics, Vol. 12, 41-197).
- Heath C., Tversky A., 1991, "Preference and Beliefs: Ambiguity and competence in choice under uncertainty", *Journal of Risk and Uncertainty* 4, 5-28.
- Hey D. J., Lotito G., Maffioletti A., 2008, "The Descriptive and Predictive Adequacy of Theories of Decision Making Under Uncertainty/Ambiguity," *University of York discussion paper in Economics*.
- Hey D. J., Orme C., 1994, "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica*, 62, 6, 1291-1326.
- Hirshleifer D., 2001, "Investor psychology and asset pricing," *Journal of Finance*, 56, 1533-1597.
- Hirshleifer D., Luo Y. G., 2001, "On the survival of overconfident traders in a competitive securities market", *Journal of Financial Markets*, 4, 73-84.
- Hirshleifer D., Shumway T., 2003, "Good day sunshine: Stock returns and the weather," *Journal of Finance*, 58, 3, 1009-1032.
- Holt C., Laury S., 2005, "Risk Aversion and Incentives: New Data without Order Effects," *American Economic Review*, 95, 3, 902-912.
- Hong H., Lim N., Stein J., 2000, "Bad News Travels Slowly: Size, Analyst Coverage, and the Profitability of Momentum Strategies", *The Journal of Finance*, 55, 265-295.
- Hong H., Stein J. C., 199 "A unified theory of underreaction, momentum trading and overreaction in asset markets", *Journal of Finance*, 54, 2143-2184.
- Hsu M., Bhatt M., Adolphs R., Tranel D., Camerer C. F., 2005, "Neural Systems Responding to Degrees of Uncertainty in Human Decision-Making," *Science*, 310, 1680-1683.
- Hvidkjaer S., 2006, "A trade-based analysis of momentum", *Review of Financial Studies*, 19, 457-491.

- Imhoff E., 1992, "The Relation between Perceived Accounting Quality and Economic Characteristics of the Firm," *Journal of Accounting and Public Policy*, 11, 2, 97-118.
- Imhoff E., Lobo G., 1984, "Information content of analysts' composite forecast revisions", *Journal of Accounting Research*, 22, 541-554.
- Jacob. J., Lys Z. T., Neale A. M., 1999, "Expertise in forecasting performance of security analysts," *Journal of Accounting and Economics*, 28, 51-82.
- Jegadeesh N., Titman S., 1993, "Returns to buying winners and selling losers: Implications for stock market efficiency", *Journal of Finance*, 48, 65-91.
- Jegadeesh N., Titman S., 2001, "Profitability of momentum strategies: An evaluation of alternative explanations," *Journal of Finance*, 56, 2, 699-720.
- Johnson T. C., 2002, "Rational momentum effects", *Journal of Finance*, 57, 585-608.
- Johnson E. J., Tversky A., 1983, "Affect, generalization, and the perception of risk," *Journal of Personality and Social Psychology*, 45, 20-31.
- Kahneman D., 1973, "Attention and effort," Englewood Cliffs, NJ: Prentice Hall.
- Kahneman D., Tversky A., 1972, "Subjective Probability: A judgment of representativeness," *Cognitive Psychology*, 3, 430-454.
- Kahneman D., Tversky A., 1973, "On the Psychology of Prediction," *Psychological Review*, 80, 237-251.
- Kahneman D., Tversky A., 1979, "Prospect Theory: An analysis of decision under risk," *Econometrica*, 47, 263-292.
- Kahneman D., Slovic P., Tversky A., 1982, *Judgement under uncertainty: Heuristics and biases*. Cambridge University Press. Cambridge UK.
- Kandel E., Pearson N., 1995, "Differential interpretation of public signals and trade in speculative markets", *Journal of Political Economy*, 103 (4), 831-872.
- Kadane J. B., Winkler R. L., 1988, "Separating probability elicitation from utilities", *Journal of American Statistical Association*, 83, 357-363.
- Kadiyala P., Rau R., 2004, "Investors reaction to corporate events announcements: Underreaction or Overreaction?" *Journal of Business*, 77, 357-386.
- Kaplanski, G., Levy, Haim, 2009, "Sentiment and stock prices: The case of aviation disasters," forthcoming, *Journal of Financial Economics*.

- Keren G., Gerritsen E. M. L., 1999, "On the Robustness and Possible Accounts of Ambiguity Aversion," *Acta Psychologica*, 103, 149-172.
- Keynes, John Maynard (1921), *A Treatise on Probability* second ed. (1948), McMillan, London.
- Keynes J. M., 1936, *The general theory of employment, interest and money*. McMillan University Press.
- Kidd J., 1970, "The utilization of subjective probabilities in production planning", 34 *Acta Psychologica*, 338.
- Klos A., Weber U. E., Weber M., 2005, "Investment Decisions and Time Horizon: Risk Perception and Risk Behavior in Repeated Gambles," *Management Science*, 51, 1777-1790.
- Klibanoff P., Marinacci M., Mukerji S., 2005, "A smooth model of decision making under ambiguity", *Econometrica*, 73, 1849-1892.
- Klibanoff P., Lamont O., Wizman T., 1998, "Investor Reaction to Salient News in Closed-End Country Funds", *Journal of Finance*, Vol. 53 (2), 673-699.
- Knight, F.H. (1921) *Risk, Uncertainty, and Profit*. Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Company.
- Kothari S. P., 2001, "Capital markets research in accounting", *Journal of Accounting and Economics*, 31, 105-231.
- Kothari S.P., Warner J. B., 1997, "Measuring long-horizon security performance", *Journal of Financial Economics*, 43, 301-339.
- Koonce L. L., McAnally L. M., Mercer M., 2000, "How do investors judge the risk of financial and derivative instruments?," *Accounting Review*.
- Korajczyk A. R., Sadka R., 2004, "Are momentum profits robust to transaction costs?," *Journal of Finance*, Vol. LIX, 3, 1039-1082.
- Kyle A. Wang A., 1997, "Speculation duopoly with agreement to disagree: Can overconfidence survive the market test?," *Journal of Finance*, 52, 5, 2073-2090.
- Lakonishok, J., Lee, I., Pearson, N., Poteshman, A., 2007, "Option Market Activity", *Review of Financial Studies*, 20, 813-857.
- Lakonishok J., Shleifer A., Vishny R., 1994, "Contrarian investment, extrapolation and risk", *Journal of Finance*, 49, 1541-1578.
- La Porta R., 1996, "Expectations and the Cross section of stock returns", *Journal of Finance*, 51(5), 1715-1742.

- La Porta R., Lakonishok J., Shleifer A., Vishny R., 1997, "Good news for value stocks: Further evidence on market efficiency," *Journal of Finance*, 52, 2, 859-874.
- Lang M., Lundholm R., 1993, "Cross-Sectional Determinants of Analyst Ratings of Corporate Disclosures," *Journal of Accounting Research*, 31,2, 246-271.
- Lee, C., Swaminathan, B., 2000, "Price momentum and trading volume," *Journal of Finance*, 55, 2017-2069.
- Leippold M., Trojani F., Vanini P., 2007, "Learning and asset prices under ambiguous information", *Review of Financial Studies*, 21, 2565-2597.
- Lee M. C. C., Shleifer A., Thaler R., "Investor sentiment and the closed-end fund puzzle," *Journal of Finance*, 46, 1, 75-109.
- Lemmon M., Portniaguina E., 2006, "Consumer confidence and asset prices: Some empirical evidence," *Review of Financial Studies*, 19, 4, 1499 – 1529.
- Livnat J., Mendenhall R., 2006, "Comparing the post-earnings announcement drift for surprises calculated from analyst and time series forecasts", *Journal of Accounting Research*, 44, 177-205.
- Livnat J., Petrovits C., 2008, "Investor sentiment, post-earnings announcement drift, and accruals," Working paper.
- Liu J., Peleg E., Subrahmanyam A., 2009, "Information, Expected Utility, and Portfolio Choice," *UCLA working paper*.
- Livnat J., Mendenhall R. R., 2006, "Comparing the post-earnings announcement drift for surprises calculated from analyst and time series forecasts," *Journal of Accounting Research*, 44, 177-205.
- Lo W. A., MacKinlay A. G., 1988, "Stock prices do not follow random walks: Evidence from a simple specifications test," *Review of Financial Studies*, 1, 41-66.
- Loewenstein G. F., Weber E. U., Hsee C. K., Welch N., 2001, "Risk as feelings," *Psychological Bulletin*, 127, 267-286.
- Lou D., 2009, "A flow-based explanation of return predictability," Yale University working paper.
- Loughran T., Ritter J. R., 1995, "The new issues puzzle", *Journal of Finance*, 50, 23-51.
- Lucas R., 1978, "Asset prices in an exchange economy", *Econometrica*, 46, 1429-1446.
- Ludvigson C. S., 2004, "Consumer confidence and consumer spending," *Journal of Economic Perspectives*, 18, 2, 29-50.

- Macceroni F., Marinacci M., Rustichini A., 2006, "Ambiguity aversion and the variational representation of preferences", *Econometrica*, 74, 1447-1498.
- Machina M J., Schmeidler D., 1992, "A More Robust Definition of Subjective Probability," *Econometrica*, 60, 4, 745-780.
- Maenhout P. J., 2002, "Robust portfolio rules and asset pricing", *Review of Financial Studies*, 17, 951-983.
- Malmendier U., Shanthikumar D., 2007, "Are small investors naïve about incentives?", *Journal of Financial Economics*, 85, 457-489.
- Matsumoto A. D., 2002, "Management's Incentives to Avoid Negative Earnings Surprises", *The Accounting Review*, 77, 3, 483-514.
- McLean, D., Zhao, M, 2009, "Investor sentiment and real investment", Working Paper, University of Alberta.
- Mehra R., Prescott E. C., "The equity premium: A puzzle", *Journal of Monetary Economics*, 15, 145-161.
- Mikhail B. M., Walther R. B., Willis H. R., 2003, "The effect of experience on security analyst underreaction," *Journal of Accounting and Economics*, 35, 101-116.
- Mikhail B. M., Walther R. B., Willis H. R., 2004, "Do Security analysts exhibit persistent differences in stock picking ability?," *Journal of Financial Economics*, 74, 67-91.
- Moskowitz J. T., Grinblatt M., 1999, "Do industries explain momentum?," *Journal of Finance*, 54, 4, 1249-1290.
- Moskowitz T. J., Vissing-Jorgensen A., 2002, "The returns to entrepreneurial investment: A private equity premium puzzle?," *American Economic Review*, 92, 754-778
- Nagel S., 2005, "Short sales, institutional investors, and the cross section of stock returns," *Journal of Financial Economics*, 78, 2, 277-309.
- O'Brien, P. C., 1987, "Individual forecasting ability," *Managerial Finance*, 13, 3-9.
- O'Brien P. C., 1990, "Forecast Accuracy of Individual Analysts in Nine Industries", *Journal of Accounting Research*, 28, 286-304.
- Odean T., 1998, "Volume, volatility, price, and profit when all traders are above average", *Journal of Finance*, 53, 1887-1934.
- Odean T., 1999, "Do investors trade too much?" *American Economic Review*, 89, 1279-1298.

- Odean T., Barber M. B., 2007, "All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors," *Review of Financial Studies*, 21, 2, 785-818.
- Oehlert G. W., 1992, "A Note on the Delta Method," *The American Statistician*, 46, 1, 27-29.
- Olsen A. R., Troughton H. G., 2000, "Are risk Premium anomalies caused by ambiguity?," *Financial Analysts Journal*, March, 24-31.
- Oskamp S., 1965, "OVERCONFIDENCE IN CASE-STUDY JUDGMENTS", *Journal of Consulting Psychology*, 29:261-5.
- Otoo, M., 1999, "Consumer sentiment and the stock market," working paper, Board of Governors of the Federal Reserve, Washington, DC.
- Quiggin J., 1992, "Generalized expected utility theory-The rank dependent model", Kluwer Academic Press, Dordrecht.
- Petersen M. A., 2008, "Estimating standard errors in finance panel data sets: Comparing approaches", *Review of Financial Studies*, 22, 435-480.
- Peng L., Xiong W., 2006, "Investor attention, overconfidence and category learning," *Journal of Finance*, 80, 3, 563-602.
- Petkova R., Zhang L., 2005, "Is value riskier than growth?," *Journal of Financial Economics*, 78, 187-202.
- Pulford B. D., 2009, "Is luck on my side? Optimism, pessimism and ambiguity aversion", *The Quarterly Journal of Experimental Psychology*, 62, 1079-1087.
- Puri, M., Robinson, D., 2007, "Optimism and economic choice," *Journal of Financial Economics*, 86, 71-99.
- Rabin M., Vayanos D., 2010, "The gambler's and hot-hand fallacies: Theory and applications", *Review of Economic Studies*, 77, 730-778.
- Read D., 2005, "Monetary incentives, what are they good for?," *Journal of Economic Methodology*, 12, 265-276.
- Rendleman et al., 1982. R.J. Rendleman, C.P. Jones and H.A. Latané , Empirical anomalies based on unexpected earnings and the importance of risk adjustments. *Journal of Financial Economics* 10, 269-287.
- Richards, R. M., 1976, "Analysts performance and the accuracy of corporate earnings forecasts," *Journal of Business*, 49, 350-357.

- Richardson S., Teoh H. S., Wysocki D. P., 2004, "The Walk-Down to Beatable Earnings Forecasts: The Role of Equity Issuance and Insider Trading Incentives", *Contemporary Accounting Research*, 21, 4, 885-924.
- Ritter, J. R., 1991, "The long run performance of Initial public offerings", *Journal of Finance*, 46, 3-27.
- Roca M., Hogarth R. M., Maule A. J., 2006, "Ambiguity seeking as a result of the status quo bias", *Journal of Risk and Uncertainty*, 32, 175-194.
- Rouwenhorst K., G., 1998, "International momentum strategies," *Journal of Finance*, 53, 267-284.
- Ross S. A., 1987, "The interrelations between finance and economics: Theoretical Perspectives", *American Economic Review*, 77, 29-34.
- Samuelson A. P., 1965, "Proof that properly anticipated prices fluctuate randomly," *Industrial Management Review*, 6, 41-49.
- Sandroni A., 2000, "Do markets favour the agents who are able to make accurate predictions?," *Econometrica*, 68, 1303-1361.
- Sarin R. K., Weber M., "Effects of Ambiguity in Market Experiments," *Management Science*, 39, 5, 602-615.
- Savage J. L., 1954, "The Foundations of Statistics," Dover Publications, New York.
- Savage L., 1971, "Elicitation of Personal Probabilities and Expectations," *Journal of the American Statistical Association*, 60, 50-62.
- Shefrin H., Statman M., 1985, "The disposition to sell winners too early and ride losers too long: Theory and Evidence", *Journal of Finance*, 40, 777-790.
- Shiller R., 1981, "Do stock prices move too much to be justified by subsequent changes in dividends?," *American Economic Review*, 71, 421-436.
- Shiller R., 2004, "Irrational Exuberance", *Princeton University Press*.
- Shleifer A., 1986, "Do demand curves for stocks slope down?," *Journal of Finance*, 41 (3), 579-590.
- Shleifer A., Vishny R., 1987, "The limits of arbitrage", *Journal of Finance*, 52, 35-55.
- Shoemaker P.J.H., 1982, "The expected utility model: its variants, purposes, evidence, and limitations", *Journal of Economic Literature*, 20, 529.

- Slovic, P., 2000, *The Perception of Risk*, Earthscan Publications, London, UK.
- Slovic P., Finucane L. M., Peters E., MacGregor G. D., 2007, "The Affect heuristic", *European Journal of Operational Research*, 177, 3, 1333-1352.
- Smith L. V., 1982, "Microeconomic Systems as an Experimental Science," *The American Economic Review*, 72, 5, 923-955.
- Sorescu S., Subrahmanyam A., 2006, "The Cross Section of Analyst Recommendations," *Journal of Financial and Quantitative Analysis*, 41, 1, 139-168.
- Statman M., Thorley S., Vorkink K., "Investor overconfidence and trading volume," *Review of Financial Studies*, 19, 1531-1565.
- Starmer C., 2000, "Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk", *Journal of Economic Literature*, 38, 332-382.
- Stickel S., 1991, "Common Stock Returns Surrounding Earnings Forecasts Revisions: More Puzzling Evidence", *The Accounting Review*, 66, 402- 416.
- Stickel S., 1992, "Reputation and Performance among Security Analysts", *Journal of Finance*, 47, 1811-1836.
- Tetlock C. P., 2007, "Giving Content to Investor Sentiment: The Role of media in the stock market", *Journal of Finance*, LXII, 3, 1139-1168.
- Thomas S., 2002, "Firm diversification and asymmetric information: Evidence from analysts' forecasts and earnings announcements", *Journal of Financial Economics*, 64, 373-396.
- Tversky A., Kahneman D., 1971, "The belief in the "law of small numbers", *Psychological Bulletin*, 76, 105-110.
- Tversky A., Kahneman D., 1974, "Judgment under uncertainty: Heuristics and biases", *Science*, 185, 1124-1131.
- Wright W. F., Bower G. H., 1992, "Mood effects on subjective probability assessment," *Organizational Behavior and Human Decision Processes*, 52, 276-291.
- Uppal R., Wang T., 2003, "Model misspecification and underdiversification", *Journal of Finance*, 58, 2465-2486.
- Wakker P., Deneffe D., 1996, "Eliciting von Neumann-Morgenstern utilities when probabilities are distorted or unknown", *Management Science*, 42, 1131-1150.
- Waymire G., "Earnings volatility and voluntary management forecast disclosure", *Journal of Accounting Research*, 23, 268-295.

Weinstein, N., 1980, "Unrealistic optimism about future life events," *Journal of Personality and Social Psychology*, 39, 806-820.

Whaley, E., R., 2009, "Understanding the VIX," *The Journal of Portfolio Management*, 35, 98-105.

Wurgler J., Zhuravskaya E., 2002, "Does Arbitrage Flatten Demand Curves for Stocks?," *The Journal of Business*, volume 75 (2002), pages 583–608.

Von Neumann J., Morgenstern O., 1944. *Theory of Games and Economic Behaviour*. Princeton University Press. Princeton. New Jersey.

Zhang F., 2006a, "Information Uncertainty and Stock Returns," *Journal of Finance*, 55, 105-136.

Zhang F., 2006b, "Information Uncertainty and Analyst Forecast Behaviour," *Contemporary Accounting Research*, 23, 2, 565-590.