

Asset Pricing and Investor Behavior

Dissertation

zur Erlangung des Doktorgrades

der Wirtschafts- und Sozialwissenschaftlichen Fakultät

der Eberhard Karls Universität Tübingen

vorgelegt von

Stephan Jank

aus München

Tübingen

2011

Tag der mündlichen Prüfung: 06.07.2012

Dekan: Professor Dr. rer. soc. Josef Schmid

1. Gutachter: Professor Dr. rer. pol. Joachim Grammig

2. Gutachter: Professor Dr. rer. pol. Christian Koziol

Acknowledgements

I am grateful to my academic teacher and supervisor Joachim Grammig for his guidance and support as well as the opportunity and freedom to explore my own research ideas. I also wish to thank Christian Koziol and Rainer Schöbel for kindly agreeing to serve on my thesis committee. Furthermore, I thank Alexander Kempf for inviting me to join the Centre for Financial Research Cologne (CFR) and for his support of my research.

Many people have accompanied me during my years as a doctoral student. Special thanks go to Michael Wedow with whom I worked on several projects at the Deutsche Bundesbank and who guided me in my early years as a researcher. I have greatly benefited from the discussions with my colleagues Thomas Dimpfl, Stefan Frey, Luis Huergo, Kerstin Kehrlé, Tobias Langen and Franziska Peter. Numerous other people have helped me by providing comments and suggestions when I presented my work at various conferences and seminars.

Last but not least, I would like to thank my parents and Marie for their invaluable support and encouragement.

Tübingen, July 2012

Stephan Jank

Contents

1	Introduction	1
2	Creative destruction and asset prices	9
2.1	Introduction	10
2.2	A simple model of creative destruction and asset prices	12
2.2.1	Technological change and asset payoffs	12
2.2.2	The household's intertemporal optimization problem	14
2.3	Data and descriptive statistics	16
2.4	Estimation results and discussion	20
2.4.1	Exposure to creative destruction risk	20
2.4.2	Model comparison	23
2.4.3	A patent activity growth-mimicking portfolio	27
2.4.4	Technological revolutions and the Fama-French factors	31
2.4.5	Robustness checks	32
2.5	Concluding remarks	33
	Appendix A: Additional Tables	35
3	Mutual fund flows, expected returns, and the real economy	39
3.1	Introduction	40
3.2	Related literature	43
3.3	Data and descriptive statistics	45
3.4	Mutual fund flows and stock market returns	51

3.5	Mutual fund flows and predictive variables	53
3.5.1	Dividend-price ratio	53
3.5.2	Other predictive variables	55
3.6	Mutual fund flows and future economic activity	63
3.6.1	Vector autoregression analysis	63
3.6.2	Forecasting comparison of market returns and fund flows	66
3.7	Concluding remarks	69
	Appendix B: Additional Tables	70
4	Can internet search queries help to predict stock market volatility?	75
4.1	Introduction	76
4.2	Data and descriptive statistics	78
4.3	The dynamics of volatility and searches	85
4.3.1	A vector autoregressive model	85
4.3.2	Do search queries add information for modeling volatility?	91
4.4	Forecast evaluation	93
4.4.1	In-sample forecast evaluation	94
4.4.2	Out-of-sample forecast evaluation	94
4.4.3	Out-of-sample forecast performance over time	100
4.5	Concluding remarks	104
5	Summary and Conclusion	105
	Bibliography	108

List of Tables

2.1	Descriptive statistics of factors	17
2.2	Descriptive statistics of portfolio excess returns	19
2.3	Time-series and cross-sectional regression	21
2.4	Expected excess return	22
2.5	Model comparison: CAPM, Fama-French and CDR model	24
2.6	Weights of the PAG-mimicking portfolio	28
2.7	Descriptive statistics: PAG-mimicking portfolio	29
2.8	PAG-mimicking portfolio and the Fama-French factors	30
A.1	Model comparison: post-war sample	35
A.2	Model comparison: equally-weighted portfolios	36
A.3	Model comparison: extended sample	36
A.4	Mimicking portfolio: time-series and cross-sectional regression	37
3.1	Summary Statistics	48
3.2	Mutual fund flows and stock market returns	52
3.3	Mutual fund flows, market returns, and changes in dividend yield	54
3.4	Testable hypotheses: Predictive variables and mutual fund flows	56
3.5	Mutual fund flows and changes in other predictive variables	59
3.6	Unexpected fund flows and changes in predictive variables	62
3.7	Mutual fund flows and economic activity	64
3.8	Mutual fund flows, market returns, and real economic activity	68

B.1	Economic activity forecasting comparison: Market return and change in dividend yield	70
B.2	Mutual fund flows and real economic activity: further specifications	71
4.1	Summary statistics	82
4.2	VAR model estimation results	88
4.3	Is search activity a helpful predictor of future volatility?	92
4.4	In-sample forecast evaluation	95
4.5	Out-of-sample forecast evaluation	98

List of Figures

2.1	Patent activity growth and the Fama-French factors	18
2.2	Fitted expected vs. realized average excess returns	26
3.1	Mutual fund flows and predictive variables	50
4.1	Realized volatility and search activity	77
4.2	Autocorrelations of realized volatility	83
4.3	Autocorrelations of search queries	84
4.4	Impulse response functions (FTSE)	90
4.5	Out-of-sample performance over time	101
4.6	Stock market volatility during the financial crisis	103

Chapter 1

Introduction

This thesis comprises three essays on empirical finance covering the topics of asset pricing and investor behavior. It examines different aspects of what determines asset prices in the long and in the short run.

Under the efficient market hypothesis the price of an asset should reflect expected discounted cash flows (see e.g. Fama 1970a). Both news about cash flows and about discount rates can affect the value of an asset. If financial markets are efficient then it follows that the expected return of an asset in excess of the risk free rate represents a risk compensation. Those assets which pay well in bad times, i.e. when investors' marginal utility is high, provide insurance for investors and thus offer lower expected excess returns. Assets paying poorly in bad times, on the other hand, have to offer a higher compensation to investors resulting in higher expected returns. The central task of empirical asset pricing is to find good proxies for investors' marginal utility.

It is important to note that the risk premia observed in financial markets are the result of an equilibrium (Cochrane 2007, 2008). This equilibrium describes the state in which each investor has settled on his or her optimal portfolio allocation. Investors prefer assets that represent an insurance and pay well in bad times. The demand for these assets drives up prices and consequently expected returns are low. In contrast, investors avoid assets which pay poorly in bad times. There is less demand for these assets, prices are lower and

expected returns are higher. Thus, it is the investors' preferences and their optimization behavior that drive asset prices up or down and generate the observed risk premia.

The efficient market hypothesis is challenged by behavioral finance (see e.g. Shleifer and Summers 1990, Barberis and Thaler 2003). Behavioral finance assumes that the demand of an asset is also driven by other non-rational factors, such as sentiment, herding or trend chasing. Furthermore, it assumes that there are limits to arbitrage. Rational investors are not able to fully arbitrage away price deviations from the fundamental value. In this setting, prices can deviate from fundamentals for longer periods of time. Much of the current research in finance centers around the question to what degree prices are determined by fundamentals, cash flow or discount rate news, and to what degree they are determined by sentiment.

In summary, the behavior of investors - rational or irrational - is important when we want to understand how asset prices evolve. This thesis looks at different aspects of the relationship between investor behavior and asset prices. Chapter 2 analyzes the cross-section of stock returns and the size and value premium. It investigates why investors prefer to hold large and growth stocks in contrast to small and value stocks. Is this preference irrational or justified by a real economic risk? Chapter 3 examines the time-varying equity premium and the portfolio adjustment of one specific investor group, namely mutual fund investors. Because mutual fund investors consist to a large extent of retail investors, flows into and out of equity funds enable us to observe the portfolio adjustment of these investors over time. The question of interest is: In which periods are retail investors willing to take more market risk and in which periods are they less willing to do so? Chapter 4 also studies the behavior of retail investors focusing on how much attention these investors pay to the stock market. It investigates whether heightened attention of retail investors, along with possible trading, contributes to stock market volatility.

In the following I will briefly introduce each chapter and describe how the results contribute to the existing literature.

The cross section of stock returns: the size and value premium

When we look at the cross section of stock returns, portfolios with low market capitalization have outperformed portfolios with a large market capitalization and portfolios with a high book-to-market value have outperformed portfolios with a low book-to-market value. This is known as the size and value premium. The higher return of small and value stocks itself is not a puzzle. However, the standard asset pricing model, the capital asset pricing model (CAPM) by Sharpe (1964), Lintner (1965) and Mossin (1966), fails to account for the differences in excess returns. That is, the differences in average excess returns of these portfolios cannot be explained by the spread in market betas across these portfolios.

In order to price size and value sorted portfolios, Fama and French (1995, 1996) augment the CAPM with two portfolios: one portfolio that is long in small stocks and short in large stocks (Small-Minus-Big, SMB) and another portfolio that is long in stocks with a high book-to-market value and short in stocks with a low book-to-market value (High-Minus-Low, HML). This Fama-French 3-factor model has become the new workhorse model in finance and is standardly used to calculate abnormal returns or outperformance. For example, HML and SMB - sometimes in addition with the momentum factor of Jegadeesh and Titman (1993) - are used to evaluate the skills of managers in mutual funds (e.g. Carhart 1997, Zheng 1999, Sapp and Tiwari 2004, Keswani and Stolin 2008). The size and value premium have also attracted attention of practitioners. While in practice the Fama-French factors HML and SMB are rarely used, it is common to evaluate the performance of a mutual fund within its peer group. This is usually done by grouping funds along size and book-to-market.

The Fama-French model has become the benchmark model in asset pricing and performance evaluation and has found its way into practice. However, there is still an ongoing debate about what kind of risk is reflected in the size and value premium. Chapter 2 “*Creative destruction and asset prices*”, which is based on joint work with Joachim Grammig, attempts to find a real economic explanation for the size and value puzzle. We argue

and provide evidence that the risk of creative destruction is priced in the size and value premium. The term “creative destruction” was coined by Schumpeter (1961) and refers to the idea that innovations make existing business models obsolete. Thus, new innovations pose a threat to existing investments. In particular, firms that are less productive are more likely to be destroyed in a technological revolution. The existing literature has identified small and value firms as firms which are less productive and which have a higher default risk (e.g. Chan and Chen 1991, Fama and French 1995, Vassalou and Xing 2004). Therefore, a technology shock has a different effect on small or large firms and on firms with low or high book-to-market value. In equilibrium, investors have to be compensated for the risk of creative destruction, which results in higher expected returns for small and value firms.

Using patent activity growth (PAG) as a proxy for technology shocks, we test whether creative destruction risk can explain the size and value premium. We find that the exposure to this factor varies along the dimensions of market capitalization and book-to-market. Returns of small and value stocks are negatively related to patent activity growth, while large and growth stocks are positively related to patent activity growth. This results in an economically significant risk premium. Since small value stocks have the highest exposure to creative destruction risk, they offer an additional 6.2 percent expected excess return per year. Large growth stocks, on the other hand, provide a hedge against creative destruction risk resulting in a discount of expected excess return of 2.4 percent annually. Overall, the creative destruction risk model can price the 25 size and book-to-market sorted portfolios and a patent activity growth-mimicking portfolio can price both HML and SMB.

Time variation of expected returns and the behavior of mutual fund investors

While the focus of Chapter 2 was on the cross-sectional variation of expected returns, Chapter 3 investigates the time variation of expected stock returns. The return of the

market portfolio over the risk-free rate, the equity premium, has been found to vary over time and the literature has identified several variables which predict the equity premium. Examples of such variables are the dividend-price ratio, interest rates or the consumption-wealth ratio (e.g. Shiller, Fischer and Friedman 1984, Fama and French 1989, Campbell 1991, Lettau and Ludvigson 2001).

The rational explanation for return predictability is that there is a time variation in risk premia. In a recession, some investors are less willing to hold risky assets and consequently will reduce their equity holdings. Those individuals who are willing to shoulder stock market risk in adverse economic times demand higher expected returns in these times. Variables that predict the equity premium have indeed a strong business cycle component. These variables are state variables that indicate bad times (e.g. recessions) or forecast these times.

Chapter 3, entitled “*Mutual fund flows, expected returns, and the real economy*”, approaches the topic of return predictability from a new angle. It looks at this asset pricing question from the perspective of the portfolio adjustment of a certain investor group, namely mutual fund investors. The group of mutual fund investors consists predominantly of private investors. Thus, the flows into and out of equity funds provide information on how private investors adjust their equity share over time. In particular, we are interested in the relation of these flows with variables that predict the equity premium.

The key finding of the analysis is that mutual fund investors seem to make just the “wrong” decisions. They sell equity when predictive variables forecast high expected returns and buy equity when predictive variables forecast low expected returns. Is this behavior irrational? Not necessarily. One has to keep in mind that *not all investors can simultaneously time the market* (Cochrane 2011). For each buyer who times the market there has to be a seller. Mutual fund investors seem to belong to the group of investors who sell equity at news of bad times and buy equity at news of good times. Different preferences or a higher exposure to labor income shocks may provide an explanation for mutual fund investors’ lower willingness to hold equity in poor economic times.

These results provide an answer to another related question: Why do mutual fund flows and stock market returns move together? There are several possible explanations for this co-movement (Warther 1995): price-pressure, feedback-trading or common response to information. The results documented in Chapter 3 are consistent with the third explanation, namely that there is a common response to macroeconomic news. I find that mutual fund flows are better described by predictive variables than by the market return alone. That is, variables that predict the economy as well as the equity premium are able to account for the positive correlation between flows into equity funds and stock market returns. Furthermore, I find that mutual fund flows are forward-looking. Fund flows predict real economic activity, which indicates that fund investors react to macroeconomic news.

Stock market volatility and retail investor behavior

Chapter 4 “*Can internet search queries help to predict stock market volatility?*”, which is based on joint work with Thomas Dimpfl, looks at retail investors’ behavior as well. We measure retail investors’ interest in the stock market by the number of internet searches for the leading stock market index in their home country. We find that search queries for stock indices rise in turbulent times, i.e. when volatility is high. If the rising interest of retail investors in the stock market triggers trading, can this influence the stock price?

The agent-based models by Lux and Marchesi (1999) and Alfarano and Lux (2007) argue that this can be the case. In these models there are two types of investors: “fundamentalists” and “noise traders”. The former follow the premise of the efficient market hypothesis, that is prices should reflect expected discounted payoffs. These investors trade if the price deviates from its fundamental value. The latter follow price trends, chart analysis or are subject to herding. Lux and Marchesi (1999) show theoretically that this behavior of noise traders can increase volatility in the market and can also generate the well-documented volatility clustering.

We find support for the agent-based models of market volatility. Granger causality between search queries and volatility is bi-directional. Heightened volatility today is followed by heightened searches tomorrow. Furthermore, increased searches today are followed by increased volatility tomorrow. Overall, retail investors' contribution to volatility is not negligible. In a long-run variance decomposition we find that log search queries account for 9% to 23% of the variance of log stock market volatility. These results are in line with recent empirical evidence by Foucault, Sraer and Thesmar (2011), who report a similar magnitude of retail investors' trading activity to the level of volatility.

The fact that search queries predict volatility is of great interest in a forecasting context, which is the main focus of Chapter 4. We utilize this finding and augment various models of realized volatility with search query data. The main results of our forecasting evaluation can be summarized as follows: Forecasting models can be significantly improved if search queries are included in the prediction equation. The improvement is evident for in-sample as well as for out-of-sample forecasts. The longer the forecast horizon, the more efficiency gains are apparent. Most importantly, search queries help to predict volatility more accurately in phases of high volatility, e.g. in the financial crisis of 2008.

Chapter 2

Creative destruction and asset prices*

Abstract

This paper introduces Schumpeter's idea of creative destruction into asset pricing. The key point of our model is that small-value firms are more likely to be destroyed during technological revolutions, while large-growth firms provide a hedge against creative destruction risk. The expected return difference between assets with the highest and lowest exposure to creative destruction risk amounts to 8.6 percent annually. A model including market return and invention growth as priced factors accounts for a large portion of the cross-sectional variation of size and book-to-market sorted portfolios and successfully prices HML and SMB.

*This chapter is based on the working paper "*Creative Destruction and Asset Prices*" by Grammig J. and S. Jank (2010).

2.1 Introduction

Historically, small stocks have outperformed large stocks and value stocks have outperformed growth stocks. These size and value premia are insufficiently explained by the Capital Asset Pricing Model (CAPM). While the Fama-French three-factor model is able to account for the size and value premia, it leaves the question of what the fundamental risk is behind HML and SMB unanswered.

This paper introduces Schumpeter's idea of creative destruction into asset pricing theory as an explanation for the size and value premia. The idea is that new and better products can render existing ones obsolete, posing an imminent risk for any investment made. This "process of industrial mutation [...] that incessantly revolutionizes the economic structure *from within*, incessantly destroying the old one, incessantly creating a new one" (Schumpeter 1961, p. 83) can be seen throughout history. Means of transportation, for example, developed within a century from horse carriage to railroad, automobile and airplane, each invention challenging the previous. Looking at the most recent technological revolution in the 1990s, inventions in the field of software and information technology led, on the one hand, to increased productivity and economic growth; on the other hand, they challenged existing business models of the music industry, media and printed newspapers. Thus, in the sense that inventions are the ultimate driver of economic growth, inventions are also the ultimate risk for an investment - namely the risk that the business idea becomes obsolete.

We propose an asset pricing model with creative destruction risk in which small and value stocks incur a higher probability of becoming destroyed during times of technological change. Previous work shows that companies with a low market value and a high book-to-market ratio are firms under distress: they are less productive and have a higher probability of default (c.f. Chan and Chen 1991, Fama and French 1995, Zhang 2005, Vassalou and Xing 2004). These distressed firms are less likely to survive technological revolutions. In equilibrium, investors have to be compensated for the risk of creative destruction, resulting in higher expected returns for small and value stocks.

Our model is a two-factor model in the spirit of Merton's (1973) Intertemporal Capital Asset Pricing Model (ICAPM). It includes market return and innovation growth, proxied by the change in patent activity as state variables. An increase of invention activity raises the risk of creative destruction and thus reduces expected cash flows of existing businesses. Long-horizon investors will prefer assets that are less exposed to creative destruction as they provide a hedge against reinvestment risk.

We find that returns of small and value stocks are negatively related to invention growth, which results in an economically significant risk premium. Small value stocks have the highest exposure to creative destruction risk and offer an additional 6.2 percent expected excess return per year. Large growth stocks, on the contrary, provide a hedge against creative destruction, resulting in a discount of expected excess return of 2.4 percent annually. The creative destruction risk model does a good job in pricing the 25 size and book-to-market sorted portfolios with the exception of the small-growth portfolio. The model is not rejected by the GMM J-test and achieves a cross-sectional R^2 of 60 percent. Finally, a patent activity growth-mimicking portfolio can price both HML and SMB, suggesting that invention growth is the real economy state variable captured by the Fama-French factors.

Our study connects several strands of literature. It relates the idea of creative destruction - an idea well established in the Schumpeterian growth theory (e.g. Segerstrom, Anant and Dinopoulos 1990, Grossman and Helpman 1991, Aghion and Howitt 1992, Helpman and Trajtenberg 1994) - to asset pricing. In this way we contribute to a growing body of literature that investigates the effects of technological innovations on asset prices (Nicholas 2008, Comin, Gertler and Santacreu 2009, Hsu 2009, Pástor and Veronesi 2009). Furthermore, we incorporate creative destruction risk into Merton's (1973) ICAPM, arguing that investment opportunities change because new technologies render existing businesses obsolete. This links our contribution to others that have empirically tested the ICAPM (e.g. Campbell 1993, 1996, Campbell and Vuolteenaho 2004, Brennan, Wang and Xia 2004).

Moreover, our work complements the literature that attempts to explain the size and

value puzzle. In particular, it refers to papers that associate market value and book-to-market ratio with measures of firm distress (e.g. Chan, Chen and Hsieh 1985, Chan and Chen 1991, Fama and French 1995). While this literature links size and book-to-market ratio to distress of *individual* firms, a connection to an *aggregate* distress factor has not been established (Lakonishok, Shleifer and Vishny 1994, Vassalou and Xing 2004). But to obtain a premium for size and value, we require a macro distress factor because idiosyncratic distress risk can be diversified away (Cochrane 2008). Our model links the individual firm's default risk to the macro variable patent activity, the proxy for creative destruction risk.

While there is evidence for a weakening or disappearing of the size premium Chen, Petkova and Zhang (2008) document that the value premium has been largely stable over time. Thus, the main challenge for asset pricing is to explain the value effect. Zhang (2005) develops a model in which costly reversibility and countercyclical price of risk generate the value premium. Petkova and Zhang (2005) show that time-varying risk goes in the right direction to explain the value premium, however it is too small to account for the observed magnitude of the value premium. For this reason they suggest considering other drivers of the risk, such as ICAPM-related risk.

2.2 A simple model of creative destruction and asset prices

2.2.1 Technological change and asset payoffs

This section presents a simple model of creative destruction that explains why small and value firms face a higher risk of being destroyed during times of technological change. The model embodies the notion that individual inventions have the potential to affect the whole economy (Aghion and Howitt 1992, Bresnahan and Trajtenberg 1995), and thus present a fundamental risk factor for investors. Examples of such pervasive inventions are the steam engine, the electric motor and the semi-conductor. Due to their impact on a wide range of sectors, Helpman and Trajtenberg (1994) refer to these inventions

as “general purpose technologies”. General purpose technologies foster productivity gains and economic growth, but they also render older technologies obsolete and destroy existing businesses. Our model explains how investors take this ambivalent nature of inventions into account, and derives implications for asset prices.

The business model of firm i generates the payoff $X_{i,t+1}$. N_t inventions occur in period t , each of which can destroy firm i with probability π_i . If π_i is small and N_t large, the number of inventions $D_{i,t+1}$ that destroy firm i follows a Poisson distribution with $\lambda_{i,t} = \pi_i \cdot N_t$. In the event that the business is destroyed ($D_{i,t+1} > 0$), the payoff $X_{i,t+1}$ equals zero. Thus, we can write the expected payoff at time t in the following way:

$$\mathbb{E}_t[X_{i,t+1}] = \exp(-N_t \cdot \pi_i) \mathbb{E}_t[X_{i,t+1} | D_{i,t+1} = 0], \quad (2.1)$$

where $P(D_{i,t+1} = 0) = \exp(-N_t \cdot \pi_i)$ gives the probability that firm i survives. The number of inventions N_t is a state variable, which influences the conditional distribution of $X_{i,t+1}$. Since more innovations have the chance of destroying the business, the expected payoff decreases when the number of inventions rises, as can be seen from

$$\frac{\partial \mathbb{E}_t[X_{i,t+1}]}{\partial N_t} = -\pi_i \cdot \exp(-N_t \cdot \pi_i) \mathbb{E}_t[X_{i,t+1} | D_{i,t+1} = 0] < 0. \quad (2.2)$$

The negative effect of an increase in inventions on the conditional expected payoff is stronger for firms with a higher individual baseline probability π_i as long as the probability that the firm survives is sufficiently high.¹ Firms with a high π_i are more exposed to the risk of destruction induced by an increase in inventions N_t .

What are the characteristics of firms with a high baseline probability of default? Vasalou and Xing (2004) provide evidence of higher default risk for value stocks. Fama and

¹Differentiating (2.2) with respect to π_i gives

$$\frac{\partial^2 \mathbb{E}_t[X_{i,t+1}]}{\partial N_t \partial \pi_i} = (\pi_i N_t - 1) \cdot \exp(-N_t \cdot \pi_i) \mathbb{E}_t[X_{i,t+1} | D_{i,t+1} = 0].$$

This expression is negative for $\lambda_{i,t} = \pi_i \cdot N_t = \mathbb{E}[D_{i,t+1}] < 1$, i.e. if the expected number of innovations that destroy the firm is less or equal to one. This corresponds to a survival probability of at least $P(D_{i,t+1} = 0) = \exp(-1) = 0.37$.

French (1995) find that value stocks are less profitable than growth stocks four years before and five years after their ranking. That small firms possess a higher default risk is shown by Chan et al. (1985) and Vassalou and Xing (2004). Furthermore, Chan and Chen (1991) find that small firms contain a large proportion of marginal firms, i.e. firms with low production efficiency. Inefficient firms may not survive times of technological change and thus face a high default risk. In summary, the previous literature identifies small and value firms as being distressed, i.e. as high π -firms.

Relating these findings to our model, it follows that the negative impact of an increase in inventions on expected payoffs should be stronger for small and value stocks. Thus, the model establishes the link between the *individual* destruction probability π_i and the *aggregate* risk factor inventions, N_t . Investors who hold stocks which are more exposed to creative destruction risk have to be compensated by higher expected returns in equilibrium.

2.2.2 The household's intertemporal optimization problem

We now outline an equilibrium model that accounts for the risk of creative destruction. The result is a two-factor model including changes in wealth and invention growth as state variables. It is a special case of Merton's (1973) ICAPM in discrete time.

In an infinite-period setting, a representative investor maximizes his or her expected life-time utility of consumption:

$$U = \mathbb{E}_t \sum_{j=0}^{\infty} \delta^j u(c_{t+j}), \quad (2.3)$$

where c_t is consumption and δ the subjective discount rate. The investor can buy a portfolio of n assets that generates wealth $W_{t+1} = R_{t+1}^W (W_t - c_t)$, where $R_{t+1}^W = \sum_{i=1}^n w_i R_i$ with portfolio weights w_i totaling one. Fama (1970b) shows that the infinite-period problem can be expressed as a two-period problem with

$$U = u(c_t) + \delta \mathbb{E}_t [V(W_{t+1}, N_{t+1})], \quad (2.4)$$

where the value function $V(\cdot)$ is defined as the maximized value of the utility function, which depends on observable state variables that account for shifts in the investment opportunity set. In our case, the value function depends on the investor's wealth W_{t+1} and the number of inventions N_{t+1} . The number of inventions captures the risk of creative destruction and the changes in investment opportunities induced by them. In a state of the world where many inventions occur - a technological revolution - it is riskier to invest in firms which are already under distress and thus might not survive. This has to be accounted for in the investor's optimization problem.

The first-order condition for optimal consumption and portfolio choice is given by

$$p_{i,t}u'(c_t) = \delta\mathbb{E}_t[V_W(W_{t+1}, N_{t+1})X_{i,t+1}], \quad (2.5)$$

where $p_{i,t}$ is the price of asset i , $X_{i,t+1}$ its payoff and $V_W(\cdot)$ refers to the derivative of the value function with respect to wealth W . Using the envelope condition $u'(c_t) = V_W(W_t, N_t)$, the stochastic discount factor can be written as

$$M_{t+1} = \delta \frac{V_W(W_{t+1}, N_{t+1})}{V_W(W_t, N_t)}. \quad (2.6)$$

First-order Taylor approximation yields the following linearized stochastic discount factor:

$$M_{t+1} = a_t + b_{1,t} \frac{W_{t+1}}{W_t} + b_{2,t} \frac{N_{t+1}}{N_t}. \quad (2.7)$$

Equation (2.5) implies the fundamental pricing equation for excess returns:

$$\mathbb{E}_t[M_{t+1}R_{i,t}^e] = 0. \quad (2.8)$$

The corresponding expected return-beta representation reads:

$$\mathbb{E}_t[R_{i,t+1}^e] = \beta_{W,t}\lambda_{W,t} + \beta_{N,t}\lambda_{N,t}, \quad (2.9)$$

where $\lambda_{W,t}$ and $\lambda_{N,t}$ capture the price of market and creative destruction risk, and $\beta_{W,t}$ and $\beta_{N,t}$ are projection coefficients which measure the asset-specific exposure to these risks.

We refer to this ICAPM with the two factors wealth portfolio and invention growth as Creative Destruction Risk (CDR) model. Note that in the case of no changes in the investment opportunity set, i.e. if the value function only depends on wealth $V(W_{t+1})$, the expected excess return of an asset is solely determined by its exposure to market risk. The model simplifies to the CAPM. But investment opportunities do change: inventions make certain businesses obsolete and create new investment opportunities. The factor invention growth, N_{t+1}/N_t , captures this change in investment opportunities. Equation (2.9) shows that an investor needs to be compensated by a higher expected return when holding assets which are more exposed to the risk of creative destruction.

2.3 Data and descriptive statistics

The key state variable in our model is invention activity. Equation (2.7) states that changes in the investment opportunity set are related to invention growth, which we approximate by the percentage change of patents issued, patent activity growth (PAG). Data on newly issued patents come from the master classification file of the United States Patent and Trademark Office (USPTO).

We argue that creative destruction risk is indeed best measured by overall patent activity growth. Of course, in hindsight some patents prove to be more relevant than others. Accounting for this difference using subsequent patent citations is an important issue when measuring the technological impact of a specific invention (Nicholas 2008). This issue loses relevance, however, when measuring creative destruction risk. *Ex-post* we observe the success or failure of an invention, and its creative destruction effects. But we are interested in the *probability* that an invention will destroy businesses. This is the risk that an investor faces *ex-ante*. We argue above that any patent has the potential to make an existing business obsolete. The example of laser technology, which revolutionized medicine, warfare, and telecommunications alike, shows the serendipitous effect of an

Table 2.1:
Descriptive statistics of factors

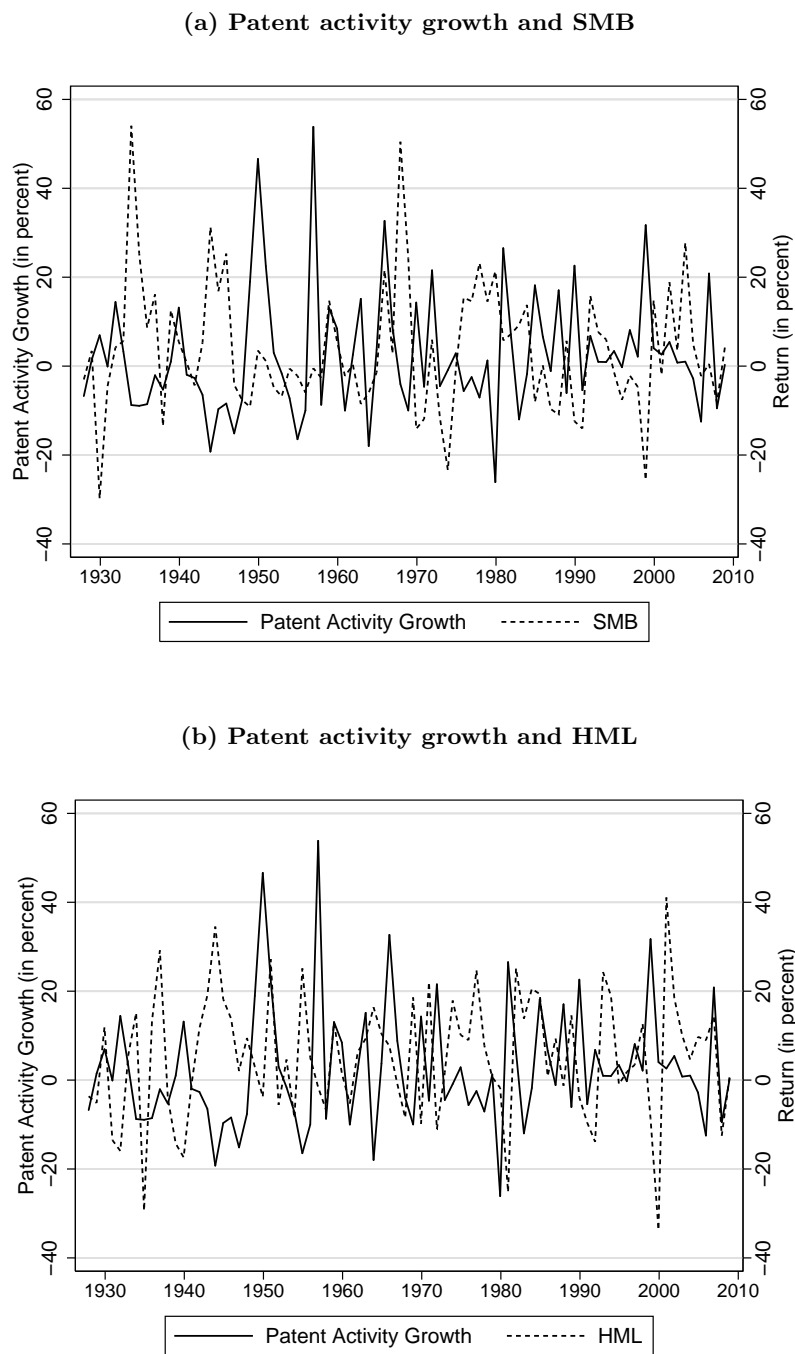
The table reports the mean, standard deviation, first-order autocorrelation AC(1) and cross-correlations of the factors market excess return (MKT), Small Minus Big (SMB), High Minus Low (HML) and patent activity growth (PAG) (all in percent). The sample period is 1927-2008, the sampling frequency is annual, and p-values are given in parentheses.

Variable	Mean	Std. Dev.	AC(1)	Correlation			
				MKT	HML	SMB	PAG
MKT	7.6	21.0	0.04 (0.71)				
HML	5.1	14.0	-0.01 (0.90)	0.11 (0.31)			
SMB	3.6	14.4	0.28 (0.01)	0.41 (0.00)	0.08 (0.50)		
PAG	2.4	13.7	0.00 (0.98)	-0.08 (0.48)	-0.21 (0.05)	-0.21 (0.06)	

innovation that was unforeseeable ex-ante (Townes 2003). It is thus the overall number of patents that best captures the risk of creative destruction.

In our main analysis we use annual data on the 25 size and book-to-market sorted portfolios ranging from 1927-2008. Data on portfolio returns and Fama-French factors are obtained from Kenneth French's homepage. We consider the longest possible sample, starting in 1927, the first available year of size and book-to-market sorted portfolios. We choose a long-run, low frequency perspective for the following reasons. First, the proxy patent activity may be prone to measurement error. The number of patents issued in a certain period can be influenced by other factors, such as institutional settings of the patent office or backlogs in the patent issuing process. These effects are presumably aggravated at higher frequencies. Furthermore, annual patent activity is arguably more suitable for capturing technological waves, which generally range over many years. The long-run perspective also complies with the ICAPM framework, in which an investor maximizes life-time utility.

Table 2.1 contains descriptive statistics on patent activity growth, market excess return and the Fama-French factors. Figure 2.1 depicts time-series of HML, SMB and patent

**Figure 2.1:****Patent activity growth and the Fama-French factors**

The graph shows patent activity growth (in percent) and the Fama-French factors Small Minus Big (SMB) and High Minus Low (HML) over the period 1927-2008.

Table 2.2:
Descriptive statistics of portfolio excess returns

The table shows summary statistics on yearly excess returns of the 25 size (vertical) and book-to-market value (horizontal) sorted portfolios from 1927-2008.

	Book-to-Market									
	Low	2	3	4	High	Low	2	3	4	High
	Mean					Standard Deviation				
Small	3.7	9.5	13.0	16.0	18.7	38.2	35.3	34.1	37.0	40.2
2	7.2	11.9	13.4	14.7	15.4	32.3	31.4	30.3	32.7	33.2
3	8.4	11.1	12.4	12.7	14.3	30.6	27.5	26.8	27.7	32.1
4	8.0	9.1	10.8	12.0	13.1	24.1	25.4	26.3	27.3	34.5
Big	7.2	7.1	8.3	8.5	10.0	21.5	19.5	22.1	25.2	31.8

activity growth. We use the value-weighted NYSE, AMEX and NASDAQ stocks as a proxy for the wealth portfolio. The market excess return (MKT) is the return of this portfolio minus the one-month Treasury Bill rate. The mean market excess return in our sample is 7.6 percent annually, which can be interpreted as the equity premium. HML (High Minus Low) is a portfolio that has long positions in stocks with high book-to-market value and short positions in stocks with low book-to-market value. Similarly, SMB (Small Minus Big) is a portfolio long in small stocks and short in large stocks.² The average premium associated with a size and value investment strategy is 3.6 percent for SMB and 5.1 percent for HML, respectively.

The size and value premia are also apparent from Table 2.2, which shows the average excess returns and standard deviations of the 25 portfolios sorted by size and book-to-market. Excess returns are computed by subtracting the one-month T-Bill rate from the raw returns. Going from left to right, value firms earn less than growth firms, and, moving from top to bottom, small firms earn more than large firms. The small-growth portfolio with an average annual excess return of just 3.7 percent is a well-known exception.

Patent activity growth averages at 2.4 percent and is considerably volatile, with a standard deviation that is comparable to HML and SMB. The PAG series shows no sign of autocorrelation and thus qualifies as a variable that captures unexpected news with

²For details on the construction of the portfolios, see Fama and French (1993).

regard to technological change. An important empirical finding, which we will elaborate on below, is that the macro variable patent activity growth is negatively correlated with both HML and SMB portfolio returns.

2.4 Estimation results and discussion

2.4.1 Exposure to creative destruction risk

Using the 25 test portfolios mentioned above, we estimate the creative destruction risk model by means of two-pass regressions and GMM, exploiting the unconditional moment restrictions implied by equation (2.8). Conditioning down and assuming time invariant parameters in (2.7), estimates of the market- and PAG-beta can be obtained by time-series regressions of excess returns on factors:

$$R_{i,t}^e = a_i + \beta_{MKT,i}MKT_t + \beta_{PAG,i}PAG_t + \varepsilon_{i,t}. \quad (2.10)$$

Factor risk premia λ_{MKT} and λ_{PAG} are estimated by a cross-sectional regression of average excess returns on beta estimates obtained in the first step. To calculate standard errors, we use the Shanken (1992) correction.

Table 2.3 displays the result of the time-series regression in Panel A. Here we report the estimates of the market beta, the patent activity growth beta and the R^2 of each time-series regression; Panel B shows the estimated factor risk premia $\hat{\lambda}_{MKT}$ and $\hat{\lambda}_{PAG}$.

The beta estimates vary considerably across portfolios with different size and book-to-market value, with a pattern that is consistent with the theoretical model of creative destruction risk. Small value firms have the strongest negative exposure to patent activity growth, with the estimate $\hat{\beta}_{PAG}$ equal to -0.42 and a t-statistic of -2.3 . Our theoretical framework suggests that these stocks possess a high baseline destruction probability π_i . A technology shock hits these firms' expected payoffs the hardest, resulting in a large drop in their prices, which corresponds to a pronounced negative beta loading.

Large growth firms, in contrast, have positive exposure to patent activity growth; the

Table 2.3:
Time-series and cross-sectional regression

Panel A contains the result of the time-series regression of excess returns on factors MKT and PAG. MKT denotes the market return in excess of the risk-free rate and PAG is patent activity growth. Test assets are the 25 portfolios sorted by size (vertical) and book-to-market value (horizontal), and the sample period is 1927-2008 at annual frequency. Beta estimates for each factor are given on the left-hand side, while t-statistics adjusted for heteroscedasticity are given on the right-hand side. The table also displays the R^2 of each regression in percent. Panel B contains the risk premia (in percent) for each factor, estimated using the cross-sectional regression of average excess returns on estimated betas. We use the Shanken (1992) correction to calculate standard errors.

Panel A: Time-Series Regression										
Book-to-Market										
	Low	2	3	4	High	Low	2	3	4	High
	$\hat{\beta}_{MKT}$					$t_{\beta_{MKT}}$				
Small	1.42	1.38	1.36	1.42	1.55	11.1	13.0	14.4	12.8	13.0
2	1.32	1.31	1.24	1.31	1.33	15.3	17.0	15.7	14.6	14.5
3	1.29	1.18	1.14	1.15	1.24	17.3	19.1	18.5	17.0	12.8
4	1.06	1.09	1.14	1.12	1.37	21.5	18.9	20.0	16.1	13.6
Big	0.97	0.89	0.97	1.06	1.28	25.7	31.1	20.7	17.3	14.2
	$\hat{\beta}_{PAG}$					$t_{\beta_{PAG}}$				
Small	-0.15	-0.24	-0.30	-0.39	-0.42	-0.78	-1.47	-2.04	-2.31	-2.31
2	-0.14	-0.18	-0.26	-0.26	-0.26	-1.03	-1.54	-2.20	-1.88	-1.86
3	-0.04	-0.20	-0.18	-0.26	-0.24	-0.35	-2.11	-1.95	-2.56	-1.60
4	0.10	-0.11	-0.14	-0.23	-0.12	1.27	-1.24	-1.55	-2.19	-0.77
Big	0.16	-0.02	-0.03	-0.08	-0.11	2.82	-0.51	-0.35	-0.84	-0.79
	R^2									
Small	61.5	68.9	73.4	69.0	69.5					
2	75.2	78.9	76.6	73.9	73.4					
3	79.2	82.7	81.8	79.5	68.5					
4	85.5	82.2	83.8	77.5	70.4					
Big	89.3	92.5	84.5	79.3	72.3					
Panel B: Cross-Sectional Regression										
	$\hat{\lambda}_{MKT}$				7.0	$t_{\lambda_{MKT}}$				2.01
	$\hat{\lambda}_{PAG}$				-14.6	$t_{\lambda_{PAG}}$				-2.06

Table 2.4:
Expected excess return

The table shows estimated expected excess returns in percent that are associated with market risk $\hat{\beta}_{MKT} \cdot \hat{\lambda}_{MKT}$ and with creative destruction risk $\hat{\beta}_{PAG} \cdot \hat{\lambda}_{PAG}$. MKT denotes market excess return and PAG patent activity growth. Estimates are taken from Table 2.3.

		Book-to-Market				
		Low	2	3	4	High
		$\hat{\beta}_{MKT} \cdot \hat{\lambda}_{MKT}$				
Small		9.9	9.6	9.5	9.9	10.8
2		9.2	9.2	8.7	9.2	9.3
3		9.0	8.2	8.0	8.0	8.7
4		7.4	7.6	7.9	7.8	9.6
Big		6.8	6.2	6.8	7.4	8.9
		$\hat{\beta}_{PAG} \cdot \hat{\lambda}_{PAG}$				
Small		2.2	3.5	4.3	5.7	6.2
2		2.0	2.7	3.9	3.8	3.8
3		0.6	2.9	2.7	3.9	3.5
4		-1.4	1.6	2.0	3.4	1.7
Big		-2.4	0.3	0.4	1.2	1.6

coefficient estimate $\hat{\beta}_{PAG}$ equals 0.16, while the t-statistic is 2.8. These stocks can generally be characterized by strong earnings growth and high profitability ratios and thus are most likely to persist throughout the technological revolution. Relatively speaking, large growth stocks might even profit from the weakness of their competitors and gain market power. This fact results in a positive beta loading with patent activity growth.

Creative destruction entails a considerable risk that is priced by the stock market. Panel B in Table 2.3 provides the $\hat{\lambda}$ estimates, which amount to 7.0 percent for the market factor and -14.6 percent for patent activity growth, significant from both a statistical and an economic point of view. Table 2.4 displays the estimated premia attributed to market risk $\hat{\lambda}_{MKT} \cdot \hat{\beta}_{MKT}$ and to creative destruction risk $\hat{\lambda}_{PAG} \cdot \hat{\beta}_{PAG}$, respectively. When we look at risk premium associated with creative destruction, small value firms earn an additional expected excess return of 6.2 percent annually due to their high risk of becoming obsolete during times of technological change. The opposite is the case for large growth firms, whose positive loading with patent activity growth leads to a discount in expected excess

returns of 2.4 percent. Overall, this yields a spread in expected excess returns of 8.6 percentage points between assets with the highest and assets with the lowest exposure to creative destruction risk.

2.4.2 Model comparison

We now compare the empirical performance of the Creative Destruction Risk (CDR) model to the CAPM (Sharpe 1964, Lintner 1965, Mossin 1966) and the Fama-French (1995) three-factor model. The CAPM can be seen as a special case of the CDR model in which investment opportunities do not change. The Fama-French model with the SMB and HML factors represents the natural benchmark for the 25 size and book-to-market sorted portfolios. The purpose of this section is not to run a horse race between the portfolio-based Fama-French model and our macro factor model. As pointed out by Cochrane (2008), portfolio-based models will have a head start on the 25 portfolios, which exhibit a correlation structure that is well captured by three principal components (see also Lewellen, Nagel and Shanken 2010). The CAPM and the Fama-French model rather serve as upper and lower benchmarks to gauge the ability of the CDR model to account for size and value premia.

GMM estimation based on the stochastic discount factor representation (2.8) provides a convenient framework for model comparisons. The stochastic discount factors M_{t+1} for CAPM, Fama-French model and CDR model are given by

$$b_0 + b_{MKT}MKT_{t+1} \quad (\text{CAPM})$$

$$b_0 + b_{MKT}MKT_{t+1} + b_{HML}HML_{t+1} + b_{SMB}SMB_{t+1} \quad (\text{Fama-French model})$$

$$b_0 + b_{MKT}MKT_{t+1} + b_{PAG}PAG_{t+1} \quad (\text{CDR model}).$$

Since we use excess returns as test assets, we de-mean all factors and set $b_0 = 1$ to ensure identification.

We report first-stage GMM estimates, with the identity matrix as a pre-specified weighting matrix, and second-stage GMM estimates using an estimate of the optimal

Table 2.5:**Model comparison: CAPM, Fama-French and CDR model**

The table contains first- and second-stage GMM results of the CAPM, Fama-French and CDR models. Test assets are the 25 size and book-to-market sorted portfolios, and the sample period is 1927-2008 at annual frequency; t-values are given in parentheses. The table also reports the GMM J-statistic and associated p-value as well as the cross-sectional R^2 in percent.

	CAPM		Fama-French Model		CDR Model	
	1st Stage	2nd Stage	1st Stage	2nd Stage	1st Stage	2nd Stage
b_{MKT}	-2.02 (-5.46)	-2.92 (-7.31)	-1.10 (-1.80)	-1.94 (-3.16)	-1.18 (-1.88)	-1.32 (-2.34)
b_{HML}			-2.76 (-3.95)	-3.53 (-4.63)		
b_{SMB}			-0.80 (-0.20)	-0.17 (0.00)		
b_{PAG}					7.54 (3.68)	5.24 (2.74)
J-statistic	46.4	39.6	36.6	29.2	29.5	34.1
p-value	0.00	0.02	0.03	0.14	0.16	0.06
R^2	25.8		81.1		59.9	

weighting matrix. Our analysis focuses on first-stage GMM results. By giving every portfolio the same weight, the model is forced to explain the size and value premium (Cochrane 2005). Second-stage GMM provides more efficient estimates, but often prices rather unusual long-short combinations of portfolios, and does not allow a comparison across models (Parker and Julliard 2005). We consider second-stage GMM results as a robustness check for our results. Following Jagannathan and Wang (1996), we report the cross-sectional R^2 as an informal and intuitive measure of goodness-of-fit.³

Table 2.5 contains first- and second-stage GMM results. Estimation of the CAPM and Fama-French model delivers the familiar results. The market excess return is a relevant pricing factor, but taken alone fails to explain the size and value premia. The R^2 is low at 26 percent, and the GMM J-test rejects the CAPM on conventional significance levels. Including SMB and HML in the stochastic discount factor, the Fama-French model performs better, although SMB is not statistically significant in this sample. The R^2

³To calculate the R^2 we run a cross-sectional regression of average realized excess returns on betas including a constant, since only in this case is the decomposition in explained and residual variation sensible. See Cochrane (2008) for further discussion.

amounts to 81 percent. Nevertheless, the J-test rejects the Fama-French model on the five percent level. Second-stage coefficient estimates for both models are similar to the first-stage results.


For the CDR model we find a significant market factor with a coefficient estimate comparable in size to the Fama-French model, and a highly significant coefficient for patent activity growth. The CDR model cannot be rejected on conventional significance levels by the first-stage GMM J-test. Second-stage GMM yields qualitatively similar results. In terms of goodness of fit, the CDR model shows a clear improvement compared to the CAPM, with an R^2 of 60 percent.

For a more detailed performance evaluation, Figure 2.2 plots average realized excess returns vs. fitted expected excess returns for the CAPM, Fama-French and CDR models. A good model fit is indicated if portfolios align along the 45-degree line. Each of the 25 test assets is numbered; the first digit refers to the size quintile and the second digit to the book-to-market quintile. For example, 15 refers to the portfolio with the smallest market value and the highest book-to-market ratio.


The first graph of Figure 2.2 depicts the well-known deficiency of the CAPM in accounting for cross-sectional return differences of size and book-to-market sorted portfolios. Unsurprisingly, the Fama-French model is more successful in pricing these portfolios. The CDR model, which includes patent activity growth in addition to the market factor, considerably improves the empirical performance as well. The model is particularly effective in pricing the small value portfolios 14 and 15. Our model of creative destruction implies that small and value firms are those with the highest risk of becoming obsolete. The additional risk premium for creative destruction thus corrects the mispricing of the CAPM.

While the CDR model generally improves the pricing of the 25 test assets, it fails to account for the small return of portfolio 11. The small-growth portfolio is well-known to present a challenge to asset pricing models (c.f. Yogo 2006, Campbell and Vuolteenaho 2004). Figure 2.2 shows that this also holds true for the Fama-French model. D'Avolio (2002), Mitchell, Pulvino and Stafford (2002) and Lamont and Thaler (2003) document

(a) CAPM



(b) Fama-French Model



(c) CDR Model




Figure 2.2:
Fitted expected vs. realized average excess returns

The figures compare fitted expected vs. realized average excess returns (in percent) given by the CAPM, the Fama-French model and the CDR model. The sample period is 1927-2008; the sampling frequency is annual. The test assets are the 25 portfolios sorted by size and book-to-market value, where the first number denotes the size quintile (1 being the smallest and 5 the largest), and the second number the book-to-market quintile (1 being the lowest and 5 the highest).

limits to arbitrage due to short-sale constraints for small-growth stocks, which offers an explanation for the difficulty to price the small-growth portfolio. The limits of arbitrage argument is also consistent with our findings from the time-series regression. Table 2.3 shows a particularly low R^2 for the small-growth portfolio, indicating that this portfolio moves less with the common risk factors, which suggests the presence of market frictions.⁴

In summary, the CDR model delivers a good performance in statistical terms and can - with the exception of the small-growth portfolio - account relatively well for the cross-sectional return differences of the 25 size and book-to-market value sorted portfolios.

2.4.3 A patent activity growth-mimicking portfolio

Can patent activity growth capture the pricing information contained in the Fama-French factors? To answer this question, we adopt a factor-mimicking portfolio approach (Breedon, Gibbons and Litzenberger 1989), acknowledging that patent activity growth may be an imperfect proxy for technological change. As pointed out by Cochrane (2008), for any macro factor that prices assets we can also use its factor-mimicking portfolio. It will contain the same pricing information, it will be less prone to measurement error, and the pricing factor will be conveniently expressed in terms of portfolio returns.

To construct the PAG-mimicking portfolio, we run the following regression:

$$PAG_t = \gamma_0 + \sum_{i=1}^K \gamma_i R_{i,t}^e + \varepsilon_t, \quad (2.11)$$

where $R_{i,t}^e$ are returns in excess of the risk-free rate of K base assets. Following Vassalou (2003), we use as base assets the six portfolios formed on size and book-to-market, which are also used to construct the Fama-French factors (for details see Fama and French 1993). Using the estimated gamma-coefficients as weights, we can form the maximum correlation

⁴The high R^2 of the Fama-French Model for all 25 portfolios in the time-series regression (c. f. Table 1, Fama and French 1996) might be a result of the inclusion of the small-growth portfolio in the construction of the SMB and HML factors.

Table 2.6:
Weights of the PAG-mimicking portfolio

The table shows the results of a time-series regression $PAG_t = \gamma_0 + \sum_{i=1}^N \gamma_i R_{i,t}^e + \varepsilon_t$ used to estimate the weights of the PAG-mimicking portfolio. Base assets are the six portfolios sorted by size and book-to-market (small-growth, small-neutral, small-value, big-growth, big-neutral and big-value (Fama and French 1993)). The sample period is 1927-2008 at annual frequency. Coefficient estimates are reported on the left-hand side; t-values are reported on the right-hand side. The table also displays the coefficient of determination R^2 (in percent) as well as the F-statistic for the hypothesis $\gamma_1 = \gamma_2 = \dots = \gamma_6 = 0$ and the corresponding p-value.

Coefficients on Base Portfolios					t-values			
	Growth	Neutral	Value	Sum		Growth	Neutral	Value
Small	0.10	-0.24	-0.09	-0.24	Small	1.18	-1.32	-0.61
Big	0.24	-0.10	0.09	0.24	Big	1.84	-0.45	0.56
Sum	0.34	-0.34	0.00		R^2		10.3	
					F-statistic		2.42	
					p-value		0.03	

portfolio that mimics the patent activity growth:

$$PAGM_t = \sum_{i=1}^K \hat{\gamma}_i R_{i,t}^e. \quad (2.12)$$

Since the base assets are zero-investment portfolios, PAGM itself is a zero-investment portfolio, and we do not require the portfolio weights to add up to one.

The estimated weights $\hat{\gamma}_i$ resulting from the time-series regression can be found in Table 2.6. As in Vassalou (2003), individual t-statistics are small due to multicollinear portfolio returns, but the estimated weights are jointly significant, as indicated by the F-test. While the presence of multicollinearity requires caution when interpreting the estimated weights (Lamont 2001), their pattern is still worth mentioning. The PAG-mimicking portfolio has long positions in value and large stocks and short positions in growth and small stocks, rather the opposite of the HML and SMB. The mimicking portfolio has maximum (positive) correlation with patent activity growth, and is thus essentially a hedge against creative destruction risk.

Table 2.7:
Descriptive statistics: PAG-mimicking portfolio

The table provides descriptive statistics for the PAG-mimicking portfolio. It displays the mean excess return, the t-value for the null hypothesis that the average excess return is equal to zero, the portfolio's standard deviation and its correlation with the market excess return (MKT), and the Fama-French factors HML and SMB. The sample period is 1927-2008 at annual frequency.

Mean		-1.66
t-value		-3.40
Std. Dev.		4.41
Correlation with:	MKT	-0.21
	HML	-0.67
	SMB	-0.66

Further properties of the PAG-mimicking portfolio are shown in Table 2.7. Its mean excess return is negative and statistically significant. The negative excess return is consistent with the idea that the PAG-mimicking portfolio is a hedge against the risk of creative destruction. Further, the mimicking portfolio shows a strong negative correlation with the Fama-French factors, implying that the PAG-mimicking portfolio explains a large proportion of the variation in these factors.

However, a pricing factor does not have to explain *all* variation in the Fama-French factors to be able to price assets comparably well. HML and SMB are neither derived from theory nor constructed to account for a specific economic risk. Only a part of HML and SMB may actually be relevant for the pricing of assets (Vassalou 2003, Petkova 2006).

To assess the pricing properties of the PAG-mimicking portfolio, we follow Cochrane (2008), who argues that macro models like the CDR model should focus on pricing the Fama-French factors rather than 25 highly correlated portfolios. Consequently, we run the following time-series regressions:

$$SMB_t = \alpha_S + \beta_{1,S}MKT_t + \beta_{2,S}PAGM_t + \varepsilon_{S,t} \quad (2.13)$$

$$HML_t = \alpha_H + \beta_{1,H}MKT_t + \beta_{2,H}PAGM_t + \varepsilon_{H,t}. \quad (2.14)$$

Since the right- and left-hand side variables of these equations are excess returns, testing

Table 2.8:
PAG-mimicking portfolio and the Fama-French factors

The table shows the results of a time-series regression of the Fama-French factors SMB and HML on the market excess return (MKT) and the patent activity growth-mimicking portfolio (PAGM). The sample period is 1927-2008 at annual frequency. α is the intercept of the time-series regression and represents the average pricing error. The table also reports the adjusted R^2 (in percent); t-values are given in parentheses.

	SMB			HML		
MKT	0.28 (3.40)		0.20 (3.59)	0.08 (1.07)		-0.02 (-0.31)
PAGM		-2.16 (-6.28)	-1.96 (-7.53)		-2.12 (-7.80)	-2.14 (-7.86)
Constant: α	1.41 (0.98)	-0.01 (-0.01)	-1.19 (-0.96)	4.56 (2.81)	1.62 (1.25)	1.73 (1.33)
Adj. R^2	16.0	43.1	50.5	0.1	43.8	43.2

for the significance of the estimated regression intercepts (i.e. pricing errors) is a test of whether the market factor and the PAG-mimicking portfolio can price SMB and HML. This is ultimately a test of whether the Fama-French factors contain additional information relevant for pricing assets.

Estimation results of the regressions (2.13) and (2.14), along with restricted versions including only MKT or PAGM as regressors, are reported in Table 2.8. Looking at SMB results, we see that the market factor prices the SMB portfolio relatively well. The beta-coefficient on MKT is significant and the pricing error is not significantly different from zero.⁵ Including PAGM in the regression, we obtain a highly significant beta estimate, the pricing error is further reduced, and the R^2 increases from 16 to 51 percent. The pricing error is actually smallest when only the PAG-mimicking portfolio is included as a regressor.

The value puzzle is reflected in the result that the market factor alone fails to price HML. The market beta is insignificant, and the pricing error of 4.5 percent is almost as large as the average return on the HML portfolio, which equals 5.1 percent (see Table 2.1).

⁵The reasonable performance of the market factor in pricing the size premium is documented by e.g. Cochrane (1999).

Once we include the factor-mimicking portfolio, we obtain a highly significant PAGM-beta, and the adjusted R^2 increases from virtually zero to 43 percent. Most importantly, the pricing error is statistically insignificant and, with only 1.7 percent, small in economic terms.

In summary, the PAG-mimicking portfolio represents a hedge portfolio against creative destruction risk and captures well the pricing information of the Fama-French factors SMB and HML.

2.4.4 Technological revolutions and the Fama-French factors

The economic rationale behind the CDR model is that cross-sectional return differences are caused by the fact that investors want to hedge creative destruction risk. This risk changes over time, which should also be reflected in stock return movements. Figure 2.1 shows that positive patent activity shocks tend to be accompanied by low returns of both HML and SMB, while negative patent activity shocks coincide with high HML and SMB returns.

We observe peaks in patent activity growth in the 1950s and 1960s, as well as the late 1990s. In the 1950s and 1960s important inventions in the field of electronics, petrochemicals and aviation were made. Computer software, digital networks and information technology were revolutionized in the 1990s. Both technology waves changed the way the economy works substantially and thus brought about creative destruction. Since small and value firms possess a higher risk of becoming obsolete during technological revolutions, prices of these assets decrease. SMB and HML returns are low. Conversely, times of low risk of creative destruction, such as the 1940s or 1970s, result in high SMB and HML returns.

Looking at the technological waves of the last century it becomes clear why they presented a substantial risk to a long-horizon investor. Consider someone who was born in 1940, started to work at the age of 20, and subsequently started investing. This would have been right in the middle of the technological revolution of the 1950s and 1960s. Assuming

a retiring age of 65, the investor would have started to consume savings in 2005, just after the peak of the information technology wave. At this point, the investor would still have had a life expectancy of 19 years.⁶ During his or her course of life, many inventions have been made, and many businesses have been destroyed.

Technology shocks were a considerable risk for this investor in the past, and still are in the retirement years to come. Large growth firms reflect efficiency, which makes them more resilient to technological shocks, providing the investor with a hedge against creative destruction risk. Small value firms, which, due to their inefficiency, are less likely to survive technological change, expose the investor to the risk of creative destruction - a risk for which the investor demands compensation.

2.4.5 Robustness checks

The results discussed in the previous sections are robust to a number of modifications. First, we confine the analysis to a post-war sample. As discussed before, our study takes a long-run, low frequency perspective in order to capture technological waves and account for the life-time horizon of the investor. The majority of empirical tests of asset pricing models, however, are conducted using post-war data sampled at quarterly frequencies. To make our results comparable, and to show that the Great Depression and the Second World War are not the main events that drive our results, we re-run the model comparison using quarterly data from 1950:Q1-2008:Q4. Table A.1 shows the results. The poor performance of the CAPM is even more severe in this period, with a cross-sectional R^2 of only 7 percent. As before, the Fama-French model achieves a high R^2 of 79 percent, but is rejected by the J-test at the 5 percent level. The results for the CDR model are confirmed: patent activity growth is a significant factor that helps to price size and book-to-market sorted portfolios, and the model is not rejected by the first-stage GMM J-test, achieving an R^2 of 56 percent, comparable in size to the long-run sample. We conclude that the Great

⁶Total population life expectancy in the United States, 2005. Source: National Vital Statistics Reports, Vol. 58, No. 10, March 3, 2010.

Depression and the Second World War do not affect our findings with regard to the role of creative destruction risk in asset pricing.

Second, we also consider a slightly different set of test assets using equally weighted portfolios. Our results are also robust for this set of test assets, as can be seen from Table A.2. The CDR model is not rejected by the J-test and the R^2 is even closer to that of the Fama-French model.

Third, we acknowledge recent criticism put forth by Lewellen et al. (2010) about the widespread use of size and book-to-market sorted portfolios in empirical asset pricing. To account for the presence of strong commonalities in these portfolios, we extend our set of test assets by ten industry portfolios. Results based on this broader sample can be found in Table A.3. Again, the results are robust in terms of parameter significance, specification test, and goodness-of-fit, and confirm the conclusions drawn from the main sample.

Finally, we compare the magnitude of the cross-sectional slopes to the factors' expected excess returns as suggested by Lewellen et al. (2010). In order to do so we repeat the time-series and cross-sectional regression using the PAG-mimicking portfolio. Results of this analysis are displayed in Table A.4. The pattern of betas is similar to the one displayed in the original analysis of Table 2.3. Since the mimicking portfolio approach reduces measurement error (c.p. Cochrane 2008), significance of β and λ estimates is even more pronounced. The estimate of the price of creative destruction risk λ_{PAGM} yields -2.1 , which is not significantly different from average excess return of the mimicking portfolio (-1.7).

2.5 Concluding remarks

This paper proposes a model of creative destruction and asset prices as an explanation for the size and value premia. Small and value firms have been shown to be under distress: they are less productive and have a higher default risk. These firms are less likely to survive technological revolutions, which results in higher expected returns for these stocks. An investor who maximizes life-time utility wants to hedge the reinvestment risk caused by

technology shocks. Hence, patent activity growth, which reflects creative destruction risk, becomes an important state variable for the investor.

The creative destruction risk model is consistent with several findings relating to the size and value effects. It is in line with the view that HML and SMB are measures of distress (e.g. Chan et al. 1985, Chan and Chen 1991, Fama and French 1995, Vassalou and Xing 2004). Further, our results are in accordance with recent findings on the Fama-French factors by Liew and Vassalou (2000) and Vassalou (2003), who show that HML and SMB forecast GDP growth. If, as we argue, technological change is the driving force behind the Fama-French factors, it should also result in greater productivity and thus higher GDP growth in the future. The same technological change that generates growth challenges existing businesses and is thus reflected in the size and value premia.

Concluding his article on efficient markets, Fama (1991) writes: “In the end, I think we can hope for a coherent story that (1) relates the cross-section properties of expected returns to the variation of expected returns through time, and (2) relates the behavior of expected returns to the real economy in a rather detailed way” (p. 1610). This paper provides such a coherent story for the size and value effect, by explaining the variation of HML and SMB through time, and linking expected returns to a fundamental risk in the real economy: the risk of creative destruction.

Appendix A: Additional Tables

Table A.1:
Model comparison: post-war sample

The table contains first- and second-stage GMM results of the CAPM, Fama-French and CDR models. Test assets are the 25 size and book-to-market sorted portfolios and the sample period is 1950:Q1-2008:Q4 at quarterly frequency; t-values are given in parentheses using Newey-West standard errors with 2 lags. The table also reports the GMM J-statistic and associated p-value as well as the cross-sectional R^2 in percent.

	CAPM		Fama-French Model		CDR Model	
	1st Stage	2nd Stage	1st Stage	2nd Stage	1st Stage	2nd Stage
b_{MKT}	-3.01 (-3.64)	-3.30 (-3.95)	-3.78 (-3.74)	-4.46 (-4.37)	-1.80 (-1.78)	-2.18 (-2.59)
b_{HML}			-6.39 (-5.64)	-6.92 (-6.09)		
b_{SMB}			-0.18 (-0.13)	0.60 (0.44)		
b_{PAG}					9.42 (3.91)	3.25 (2.03)
J-Statistic	41.6	41.3	35.4	34.7	30.8	38.3
p-value	0.01	0.02	0.04	0.04	0.13	0.02
R^2	6.6		78.9		56.1	

Table A.2:**Model comparison: equally-weighted portfolios**

The table contains first- and second-stage GMM results of the CAPM, Fama-French and CDR models. Test assets are the 25 size and book-to-market sorted portfolios, equally weighted. The sample period is 1927-2008 at annual frequency; t-values are given in parentheses. The table also reports the GMM J-statistic and associated p-value as well as the cross-sectional R^2 in percent.

	CAPM		Fama-French Model		CDR Model	
	1st Stage	2nd Stage	1st Stage	2nd Stage	1st Stage	2nd Stage
b_{MKT}	-2.15 (-6.65)	-2.71 (-7.16)	-0.64 (-1.01)	-1.23 (-1.90)	-0.98 (-1.55)	-0.79 (-1.34)
b_{HML}			-3.05 (-4.58)	-3.72 (-5.18)		
b_{SMB}			-1.73 (-1.18)	-1.01 (0.00)		
b_{PAG}					9.61 (5.20)	9.01 (5.07)
J-Statistic	44.6	41.2	37.9	34.2	29.4	31.2
p-value	0.01	0.02	0.02	0.05	0.17	0.12
R^2	51.2		85.4		75.4	

Table A.3:**Model comparison: extended sample****10 industry and 25 size and book-to-market sorted portfolios**

The table contains first- and second-stage GMM results of the CAPM, Fama-French and CDR models. Test assets are the 25 size and book-to-market sorted portfolios and 10 industry portfolios. The sample period is 1927-2008 at annual frequency; t-values are given in parentheses. The table also reports the GMM J-statistic and associated p-value as well as the cross-sectional R^2 in percent.

	CAPM		Fama-French Model		CDR Model	
	1st Stage	2nd Stage	1st Stage	2nd Stage	1st Stage	2nd Stage
b_{MKT}	-2.01 (-6.24)	-2.88 (-8.42)	-1.61 (-2.92)	-2.15 (-3.95)	-1.58 (-3.63)	-1.48 (-3.33)
b_{HML}			-2.16 (-3.32)	-3.01 (-4.54)		
b_{SMB}			-0.06 (0.12)	0.09 (0.00)		
b_{PAG}					4.24 (3.45)	5.72 (4.87)
J-Statistic	58.7	51.6	56.4	47.4	46.7	43.3
p-value	0.01	0.03	0.00	0.04	0.06	0.11
R^2	29.0		69.8		52.9	

Table A.4:
Mimicking portfolio: time-series and cross-sectional regression

This table repeats the time-series and cross-sectional regression of Table 2.3 using the patent activity growth mimicking portfolio (PAGM). Panel A displays the estimates of the time-series regression, Panel B the estimated price of risk λ from the cross-sectional regression. In (1) we test whether the price of risk is different from zero, in (2) whether it is different from the factors' average return.

Panel A: Time Series Regression										
Book-to-Market										
	Low	2	3	4	High	Low	2	3	4	High
	$\hat{\beta}_{MKT}$					$t_{\beta_{MKT}}$				
Small	1.38	1.29	1.24	1.27	1.39	10.8	13.5	18.6	17.9	18.0
2	1.28	1.23	1.13	1.19	1.21	15.2	19.9	22.3	21.8	19.3
3	1.28	1.11	1.06	1.05	1.12	16.9	22.8	26.2	25.8	16.0
4	1.08	1.04	1.07	1.05	1.28	21.8	20.2	25.1	18.9	14.2
Big	1.02	0.89	0.94	1.02	1.23	37.4	30.5	21.1	17.6	14.1
	$\hat{\beta}_{PAGM}$					$t_{\beta_{PAGM}}$				
Small	-1.01	-2.23	-3.09	-3.91	-4.20	-1.66	-4.88	-9.67	-11.56	-11.41
2	-1.18	-2.17	-2.77	-3.22	-3.06	-2.95	-7.40	-11.45	-12.39	-10.27
3	-0.45	-1.80	-2.11	-2.49	-3.03	-1.24	-7.79	-11.03	-12.77	-9.07
4	0.48	-1.26	-1.72	-2.03	-2.12	2.05	-5.14	-8.48	-7.69	-4.94
Big	1.26	-0.02	-0.76	-1.06	-1.31	9.74	-0.15	-3.60	-3.83	-3.16
	R^2									
Small	62.5	75.5	87.2	87.7	87.7					
2	77.4	87.2	90.7	90.7	88.1					
3	79.6	89.7	92.5	92.7	84.1					
4	85.9	86.4	91.3	86.3	77.2					
Big	94.7	92.5	86.7	82.4	75.2					
Panel B: Cross-Sectional Regression										
	$\hat{\lambda}$	t-value ⁽¹⁾			Mean return			t-value ⁽²⁾		
MKT	6.2	2.33			7.6			-0.54		
PAGM	-2.1	-3.50			-1.7			-0.74		

Chapter 3

Mutual fund flows, expected returns, and the real economy*

Abstract

This paper investigates the relation between mutual fund flows and the real economy. The findings of this paper support the theory that the positive co-movement of flows into equity funds and stock market returns is explained by a common response to macroeconomic news. Variables that predict the real economy as well as the equity premium - in particular dividend-price ratio, default spread, relative T-Bill rate and consumption-wealth ratio - are related to fund flows and can account for the correlation of flows and market returns. Furthermore, consistent with the information-response hypothesis, mutual fund flows are forward-looking and predict real economic activity.

*This chapter is based on the working paper “*Mutual fund flows, expected returns, and the real economy*” by Jank, S. (2010).

3.1 Introduction

Stock market returns and flows into equity funds are contemporaneously correlated: positive returns are accompanied by inflows into equity funds, and negative returns are accompanied by outflows or diminished inflows. Several competing theories provide explanations for this co-movement (Warther 1995). The so-called *feedback-trader hypothesis* states that market returns cause fund flows. Investors buy fund shares as a response to rising prices and sell when prices fall, hereby causing the positive co-movement. But causality could also run the opposite way. Mutual fund investors may represent sentiment unrelated to fundamentals. Through this uninformed demand by fund investors stock prices may temporarily diverge from their fundamental values. This hypothesis, which claims that flows cause returns, is known as the *price-pressure hypothesis*. A third explanation, the *information-response hypothesis*, states that both stock market returns and fund flows together react to new information.

This paper adopts a new approach in testing whether reaction to information explains the co-movement of fund flows and returns. In particular, the paper explores whether a specific sort of information, namely macroeconomic information, is related to mutual fund flows. I take two indirect methods for testing this. First, I consider predictive variables as proxies for macroeconomic news. These predictive variables are forward-looking, i.e. they predict real economic activity as well as the equity premium. If mutual fund flows react to information about the real economy we should observe a co-movement of flows and first differences of these forward-looking variables. Second, I investigate if mutual fund flows in themselves contain information. If mutual fund investors respond to information, e.g. by buying at good news and selling at bad news, and if they are *on average* right, then the state of the economy should be worse after outflows and better after inflows into mutual funds (see e.g. Roll 1984, for a similar argument). Thus, if mutual fund investors react on macroeconomic news, then mutual fund flows, along with stock market returns, should be able to predict economic activity.

The results of this paper can be summarized as follows. Mutual fund flows are in-

deed related to predictive variables, and in particular to dividend yield. In line with the information-response hypothesis, mutual fund flows are also related to other variables that predict the equity premium and the real economy: an increase in default spread or consumption-wealth ratio (*cay*), both indicating riskier times, is associated with outflows; an increase in relative T-Bill rate, indicating less risky times, is associated with inflows into equity funds. Overall, predictive variables describe fund flows considerably better than stock market returns alone. While stock market return explains about 40.8 percent of the variation of unexpected mutual fund flows, predictive variables explain up to 51.7 percent. Furthermore, predictive variables can account for the correlation between fund flows and market returns. With regard to the second hypothesis, I find that mutual fund flows - like stock prices - are forward-looking. Mutual fund flows predict future economic activity, measured by real GDP, industrial production, consumption and labor income. These findings support the theory that market returns and mutual fund flows simultaneously react to macroeconomic news.

These results are consistent with other studies that analyze aggregate fund flows and market returns (Warther 1995, Edelen and Warner 2001, Rakowski and Wang 2009). While these studies find evidence in favor for information as common driver of both flows and returns, their findings are ambiguous with respect to other explanations. Warther (1995) concludes that the co-movement of flows and returns is either explained by response to information or by price pressure. In order to disentangle competing theories Edelen and Warner (2001) turn to daily data, but despite the high frequency their results are consistent with either a common response to information or feedback trading. This paper looks at the information-response hypothesis from a different angle and develops two new testable implications, which when tested provide additional support for the information-response hypothesis. Moreover, this paper addresses the question, *which* information matters to mutual fund investors by showing that macroeconomic information is an important determinant of fund flows.

The results of this paper are not only interesting for the question of what explains

the co-movement of flows and market returns but also for the question of portfolio choice and tactical asset allocation. A wide literature explores how investors can use predictive variables in order to improve their portfolio performance (e.g. Brennan, Schwartz and Lagnado 1997, Campbell and Viceira 1999, Barberis 2000, Campbell and Viceira 2002, among others).¹ From a tactical asset allocation standpoint mutual fund investors seem to make just the “wrong” decisions: Mutual fund investors sell stocks, when predictive variable signal high expected returns, and they buy stocks when predictive variables signal low expected returns. The important thing to note, however, is that not *all* investors can follow a tactical asset allocation strategy (Cochrane 2011). Someone has to take the other side of each buy or sell transaction. And this decision of course depends on *differences* in investors and their preferences.

One can also look at this from a different perspective: It is exactly *because* some investors sell at news of bad times, that we observe time-varying expected returns in the first place. Take the rational explanation of why expected returns change over time (e.g. Fama and French 1989, Cochrane 1994, Lettau and Ludvigson 2001): in a recession, some people are less willing to hold risky assets and consequently will reduce their equity holdings. Those investors who are willing to shoulder stock market risk in adverse economic times have to be compensated in equilibrium, which results in higher expected returns in bad times. The results presented in this paper suggest that mutual fund investors belong to the group of investors who are less willing to hold equity in poor economic times.² Thus, mutual fund investors responding to macroeconomic news and an equity premium varying over the business cycle can be seen as two sides of the same coin.

Different preferences or high idiosyncratic income risk may be the reason for mutual fund investors’ lower willingness to hold equity in poor economic times. Mutual fund

¹For a summary of the literature see Cochrane (2007).

²Theoretically, there are two other cases: First, if mutual fund investors do not differ from the average investor, then bad news should lead to negative returns but no portfolio adjustment by mutual fund investors. In this case we would observe no correlation between fund flows and stock market returns. Second, if mutual fund investors tend to take more risk in bad times, then bad news should lead to negative returns and positive inflows. In this case we would observe a negative correlation between fund flows and market returns. Thus, the fact that we observe a positive correlation between fund flows and returns is consistent with the theory that mutual fund investors are reacting to information *and* that they are less than average willing to hold equity in bad times.

investors are predominantly private investors, who are probably more severely affected by a recession than their institutional counterparts. Moreover, within the group of retail investors mutual fund investors are special. Mutual funds provide a low cost access to the equity market (Fama and French 2002) allowing certain investors, which may not have done so otherwise, to participate in the stock market. These investors, however, are presumably more affected by economic contractions and thus more likely to sell stocks when there is bad news about the economy.

The findings presented in this paper also offer a new perspective on the question of the performance of mutual fund investors as a group. Nesbitt (1995), Friesen and Sapp (2007) and Ben-Rephael, Kandel and Wohl (2010) for example, give evidence that mutual fund investors have poor market timing ability - that is, they earn lower returns than the market. The results of this paper provide a simple explanation for lower returns realized by fund investors. Mutual fund investors seem to be less willing to bear risk in bad times, and therefore should also earn a lower expected return in equilibrium.

3.2 Related literature

This paper connects and contributes to several strands of literature. First and foremost, it expands the literature that investigates aggregate fund flows and their relation to stock market returns. Warther (1995), one of the first to examine fund flows and their relationship to security returns, documents a significant contemporaneous correlation between stock market returns and mutual fund flows at a monthly frequency. As regards explanations for this co-movement, Warther concludes that stock returns and fund flows move together either because of price pressure or because of a common response to information. The return-reversal tests performed by Warther provide no evidence for the presence of price pressure and thus point to the information-response hypothesis, however the reversal tests are admittedly not very powerful.

To disentangle causality between flows and returns, Edelen and Warner (2001) turn to high-frequency data. However, evidence with regard to one or the other explanation

is, despite the high frequency, mixed. This paper takes a different approach. Rather than examining high frequency flows, I investigate low frequency flows and their link to the real economy, since ultimately the decision to invest into financial assets cannot be isolated from the real economy. The results at this lower frequency are consistent with the results at higher frequencies, e.g. with Rakowski and Wang (2009) who find a dominant information effect in fund flows.

This article links the studies on aggregate mutual fund flows to the broader literature on time-varying equity premium and asset prices. Several variables have been found to predict the equity premium, and these predictive variables are related to the business cycle.³ In this paper I argue that mutual fund flows reacting to macroeconomic news and an equity premium varying over the business cycle can be seen as two sides of the same coin. The link between mutual fund flows and predictive variables provides new evidence with regard to investor heterogeneity (see, e.g., Mankiw 1986, Dumas 1989, Constantinides and Duffie 1996, Wang 1996, Grossman and Zhou 1996, Chan and Kogan 2002). The paper contributes to this literature by demonstrating that one specific group of investors, mutual fund investors, are less willing to hold equity in adverse economic times and sell at news of such times. We thereby learn which investors are willing to hold equity throughout the business cycle.

Time-varying risk premia on the stock market are closely intertwined with the real economy. Consequently, this paper is also connected to the body of literature that documents a strong relationship between stock returns and future economic activity.⁴ In particular it refers to the studies by Fama (1990) and Schwert (1990) which show that stock market returns predict future economic activity. The paper augments this literature by showing that not only stock market returns, but also mutual fund flows are forward-looking and predict real economic activity.

In a related paper Chalmers, Kaul and Phillips (2010) also document a flight-to-quality among mutual fund investors during economic crises. The results presented in this paper

³See Table 3.4 for a summary of the literature.

⁴E.g. Fama (1981), Geske and Roll (1983), Kaul (1987), Fama (1990), Schwert (1990), Barro (1990) and Chen (1991).

differ in important parts and go beyond their analysis. This paper shows that (unexpected) changes in predictive variables (proxying for macroeconomic news), rather than the level of these variables, are an important determinant of mutual fund flows. In this context I further test competing theories for the co-movement of market returns and mutual fund flows by showing that changes in predictive variables can account for the correlation of flows and returns. Moreover, this paper relates the behavior of mutual fund investors to the broader literature and theory on the time-varying equity premium.

3.3 Data and descriptive statistics

Data on aggregate flows into equity funds are provided by the Investment Company Institute (ICI).⁵ Following Warther (1995) and Fant (1999) I calculate quarterly net flows as new sales minus redemptions plus exchanges-in minus exchanges-out, and standardize flows by the total market value of the previous quarter using the total market index from Thomson Reuters Datastream. Fund flows are measured over the period of one quarter in order to link them to macroeconomic data. For example, the consumption-wealth ratio, which serves as a key predictive variable in our context, is only available at quarterly frequency. Overall, the mutual fund data cover 26 years, from 1984:Q1 until 2009:Q4.

The market return is proxied by the return of the S&P 500, which is also obtained from Thomson Reuters Datastream. Several variables that predict the equity premium are considered in this paper: dividend-price ratio, default spread, term spread, relative T-Bill rate and the consumption-wealth ratio. The dividend yield of the S&P 500 is measured by the ratio of average annual dividends and end-of-quarter prices, taken in logs. Data on dividends and prices of the S&P 500 are taken from Robert Shiller's homepage. The default spread is calculated as the end-of-period difference between Moody's BAA and AAA Seasoned Corporate Bond Yield. Term spread is computed as the difference between the 10-year and 1-year maturity Treasury rates at the end of each quarter. Following

⁵ICI data cover about 98 percent of assets in the mutual fund industry (see, e.g., ICI - Trends in Mutual Fund Investing, July 2010).

Campbell (1991) and Hodrick (1992), who use a stochastic detrended T-Bill rate to forecast returns, the relative T-Bill rate is calculated as the 3-month T-Bill rate minus its 12-month moving average. Data on corporate bonds and Treasury rates are all obtained from the FRED database of the Federal Reserve Bank of St. Louis. The updated time series of the consumption-wealth ratio cay is from Martin Lettau and Sydney Ludvigson.⁶

All the predictive variables mentioned above are measures for the state of the economy. They either relate to present economic activity or even more importantly they relate to economic activity in the near future. This paper aims to investigate whether fund flows react to news about the real economy. Such news would be reflected in a *change* in these predictive variables, and if fund flows react to such news then unexpected flows should be contemporaneously related to these shifts in predictive variables. For this reason I consider first differences (indicated by Δ) of all predictive variables when investigating their relationship to unexpected fund flows. This parsimonious way to measure unexpected changes in state variables follows the practice of Chen, Roll and Ross (1986). As can be seen from Table 3.1 first differences of predictive variables show no significant serial correlation and thus can be understood as unanticipated.

An alternative way to measure news would be to estimate a vector autoregression (VAR) model and using its residuals as innovations. However, for several reasons I follow the more parsimonious approach of taking first differences. Due to their forward-looking nature changes of predictive variables should, economically speaking, already reflect news. Furthermore, a VAR model including all predictive variables would be highly parametrized and might be less robust out of sample. Moreover, model misspecification and errors-in-variables are problematic when these variables later are used as explanatory variables. For further discussion on these issues see Chen, Roll and Ross (1986).⁷

The main measure for economic activity is real gross domestic product (GDP). In addi-

⁶The consumption-wealth ratio cay is computed as $cay_t = c_t - 0.2084a_t - 0.6711y_t$ and demeaned, where c_t is consumption, a_t asset wealth, and y_t labor income. For further details, see Lettau and Ludvigson (2001, 2004).

⁷As a robustness check I also estimate a VAR(1) model of all predictive variables (dividend-price ratio, default spread, term spread, T-Bill rate and cay) jointly with market return. Estimated innovations from this model are highly correlated with first differences of these variables.

tion, I also consider real industrial production, consumption and labor income as measures for the state of the economy. Several studies link stock returns to these macroeconomic variables, including Fama (1990), Schwert (1990), Chen (1991) and Lamont (2001). As the corporate bonds and Treasury rates data, all macroeconomic data are also from FRED database of the Federal Reserve Bank of St. Louis.

Table 3.1 exhibits descriptive statistics of mutual fund flows, stock market returns and the first difference of predictive variables, as well as GDP growth. Panel A provides mean, standard deviation and autocorrelations, while Panel B shows correlations. In order to illustrate the relation of market return, fund flows and predictive variables to the real economy, the table also reports correlations to past, future and contemporaneous GDP growth. The relation to the other measures of economic activity (industrial production, consumption and labor income growth) is similar, but not reported for reasons of brevity.

Panel B displays the correlation matrix of the variables mentioned above. First of all, there is a strong co-movement of mutual fund flows and stock market returns with a correlation coefficient of 0.46. But other variables also show a notable correlation with mutual fund flows, in particular $\Delta(d-p)_t$ and Δcay_t . Note that the correlation between fund flows and $\Delta(d-p)_t$ is even stronger than between flows and returns with a correlation of -0.53. Most of the predictive variables are correlated with each other, and they are correlated with, or predict real economic activity. In addition, market returns and mutual fund flows are also positively related to contemporaneous and future GDP growth.

Figure 3.1 illustrates the business cycle pattern of mutual fund flows and the most important predictive variables: the dividend-price ratio and the consumption-wealth ratio cay . Just before and during recessions there is a surge in both dividend yield and cay , and at the same time there are outflows from equity funds. The rise in the dividend-price ratio and cay accompanied by outflows from equity funds is particularly strong for the most recent and most severe recession.

Table 3.1:
Summary Statistics

This table reports summary statistics of net flows into equity funds (in percent) normalized by lagged total market capitalization, market returns (in percent) measured by the return of the S&P 500 index, the first difference of predictive variables and GDP growth (in percent). $\Delta(d-p)_t$ is the change in the dividend-price ratio, $\Delta\text{Default}_t$ is the change in the default spread, ΔTerm_t the change in the term spread, $\Delta\text{Rel. T-Bill}_t$ the change in the relative 3-month T-Bill rate and Δcay_t the change in the consumption-wealth ratio. Panel A provides mean, standard deviation and autocorrelations; Panel B displays correlations.

Panel A: Mean, Standard Deviation and Autocorrelations

Variable	Mean	Std. Dev.	Autocorrelations for Lag						
			1	2	3	4	5	10	
Flow _t	0.35	0.44	0.72	0.55	0.51	0.48	0.38	0.19	
Return _t	3.11	8.55	0.07	0.00	0.02	0.08	-0.02	0.04	
$\Delta(d-p)_t$	-0.01	0.08	0.13	-0.01	-0.02	0.01	-0.03	0.00	
$\Delta\text{Default}_t$	0.00	0.26	0.14	-0.26	-0.13	-0.01	-0.11	-0.01	
ΔTerm_t	0.02	0.42	0.20	0.25	0.06	-0.11	-0.01	0.03	
$\Delta\text{Rel. T-Bill}_t$	0.00	0.52	-0.05	0.13	-0.17	-0.21	0.08	-0.02	
Δcay_t	0.00	0.01	-0.19	0.00	-0.09	0.04	-0.08	-0.08	
GDP Growth _t	0.69	0.63	0.46	0.42	0.13	0.15	0.05	0.02	

Table 3.1 -Continued

Panel B: Correlation Matrix

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
(1) Flow _t	1									
(2) Return _t	0.46	1								
(3) $\Delta(d-p)_t$	-0.53	-0.92	1							
(4) $\Delta\text{Default}_t$	-0.22	-0.32	0.44	1						
(5) ΔTerm_t	-0.12	-0.09	0.06	-0.12	1					
(6) $\Delta\text{Rel. T-Bill}_t$	0.19	0.11	-0.20	-0.09	-0.47	1				
(7) Δcay_t	-0.32	-0.58	0.53	0.24	0.06	-0.16	1			
(8) GDP Growth _{t-1}	0.26	0.14	-0.16	0.04	-0.31	-0.04	0.05	1		
(9) GDP Growth _t	0.37	0.20	-0.30	-0.18	-0.22	0.17	-0.14	0.47	1	
(10) GDP Growth _{t+1}	0.44	0.32	-0.42	-0.26	-0.12	0.05	-0.25	0.44	0.47	1

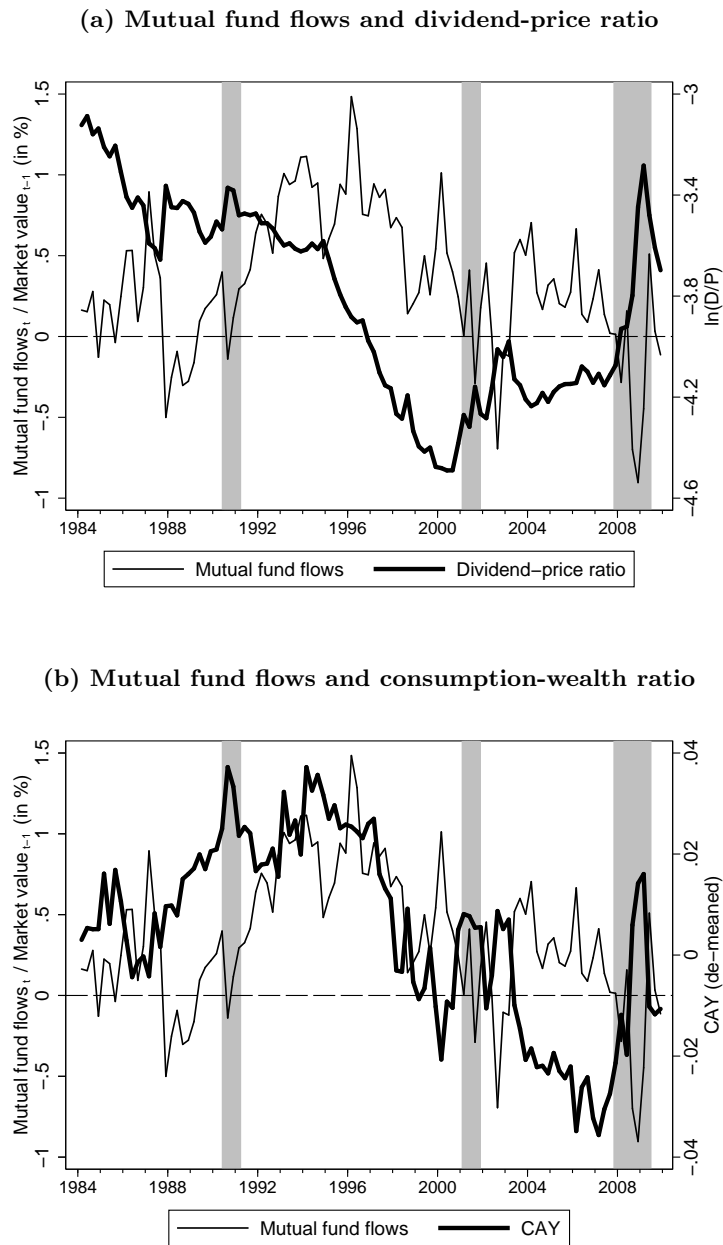


Figure 3.1:
Mutual fund flows and predictive variables

This table displays flows into equity mutual funds (in percent) normalized by dividing by lagged total market capitalization and their relation to a) the log dividend-price ratio and b) the consumption-wealth ratio *cay*. Shaded areas indicate recessions as defined by the National Bureau of Economic Research, NBER. The sample period is 1984:Q1-2009:Q4.

3.4 Mutual fund flows and stock market returns

I begin by analyzing the properties of aggregate mutual funds and their relation to stock market returns. Following Warther (1995) I run a regression of mutual fund flows on its lag and concurrent market returns, the results of which can be found in Table 3.2. In line with Warther's findings, column (1) shows that quarterly fund flows can be modeled by an AR(1)-process. The coefficient of the first lag is 0.73 and statistically significant, and the Ljung-Box Q-statistic is unable to detect any remaining autocorrelation in the residuals. Mutual fund flows show a sizable contemporaneous correlation with stock returns, as demonstrated in column (2). The share of mutual fund flow variance explained by market returns amounts to 20.8 percent. In column (3) both regressors, past flows and concurrent returns, are included and coefficient estimates are virtually the same as before.

Once more following Warther (1995), fund flows are separated into their expected and unexpected components, where the expected component is the fitted values of the AR(1)-model estimated in column (1), and the unexpected component is its residuals. Comparing columns (4) and (5), we observe that market returns are correlated with unexpected flows, but are uncorrelated with expected flows. The result of column (4) underlines the strong relation between market returns and flows, with market returns explaining up to 40.8 percent of flow innovations. This regression should not be read in a causal sense, i.e., that returns cause flows: it merely measures the linear dependence between flows and returns, and one can likewise run a regression of returns on flows. The results with respect to the R^2 , the explained variation, are the same, i.e., 40.8 percent of returns are explained by mutual fund flow innovations.

The separation of mutual fund flows into expected and unexpected flows provides a direct insight into the relation between returns and fund flows. For the remainder of the analysis, however, I will use the multiple regression model presented in column (3), since according to the Frisch-Waugh Theorem, the market return's regression coefficients of the partial model (4) and multiple regression model (3) have to be equal. This avoids any

Table 3.2:
Mutual fund flows and stock market returns

The table shows the results of a regression of net flows into equity funds on past flows and contemporaneous market returns. The R^2 (simple and adjusted) is provided for each regression. Column (1) displays the Ljung-Box Q-statistic for the test that residuals are not autocorrelated (up to lag 20). Unexpected and expected net flows in columns (4) and (5) are the residuals and fitted values of the regression model in column (1). Heteroskedasticity-robust t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5% and 1% level respectively.

	(1)	(2)	(3)	(4)	(5)
	Dependent Variables:				
	Flow _t	Flow _t	Flow _t	Unexpected Flow _t	Expected Flow _t
Flow _{t-1}	0.73*** (10.92)		0.72*** (13.52)		
Return _t		0.02*** (5.30)	0.02*** (8.41)	0.02*** (8.45)	0.00 (0.27)
Constant	0.09** (2.30)	0.27*** (6.69)	0.02 (0.80)	-0.07*** (-2.95)	0.35*** (10.24)
R^2	52.5	21.6	72.2	41.4	0.1
Adj. R^2	52.1	20.8	71.6	40.8	-0.9
Q-Statistic	20.0				
p-value	0.46				

complications regarding inference, which might arise due to the fact that unexpected flows are estimated.

Table 3.2 documents facts about fund flows and stock market returns, yet there are several explanations for this phenomenon (Warther 1995, Edwards and Zhang 1998, Fant 1999). The first explanation for the co-movement of stock market returns and equity fund flows is the so-called feedback-trader hypothesis, in which mutual fund investors react to positive returns with inflows and to negative returns with outflows. Another explanation is the price-pressure hypothesis. In this setting, mutual fund investors represent investor sentiment, i.e. optimism or pessimism unrelated to fundamentals. These uniformed trades induced from fund investors should not affect the information-efficient price in the long run, but due to higher demand prices will temporary diverge from their fundal value.

A further possibility is that information drives both returns and flows. In this case mutual fund investors react to information (possibly along with other investors) and the market price efficiently reflects this new information as well. This study focuses on the last explanation and its implications, examining in particular whether both fund flows and returns together react to a specific sort of information, namely news about the macroeconomy.

In both cases, under the price-pressure hypothesis and the information-response hypothesis, the mutual fund investors' demand for equity changes. The crucial difference between the two explanations, though, is that in case of the price-pressure hypothesis fund flows are unrelated to fundamentals, but in case of the information-response hypothesis they are driven by fundamental news about the economy.

The information-response hypothesis has two main implications, which are tested in the following. First, variables that predict the real economy should be related to mutual fund flows. Second, if mutual fund flows react to news about the real economy, then mutual fund flows should also predict real economic activity.

3.5 Mutual fund flows and predictive variables

3.5.1 Dividend-price ratio

To test the information-response hypothesis I first explore the connection between mutual fund flows and changes in the dividend yield. The dividend-price ratio or dividend yield is one of the most common variables used to predict the equity premium (see, e.g., Shiller, Fischer and Friedman 1984, Fama and French 1988, Campbell and Shiller 1988, Ferson and Harvey 1991). A high dividend-price ratio forecasts a high market excess return. In riskier times prices are low in relation to dividends, and the dividend-price ratio is high. During these times investors are less willing to hold equity, and those investors who are willing to hold equity need to be compensated by a higher expected return.

News about riskier or less risky times is thus captured by a *change* in the dividend-

Table 3.3:
Mutual fund flows, market returns, and changes in dividend yield

The table shows the results of a regression of net flows into equity funds on past flows, contemporaneous market returns and changes in dividend-price ratio. The R^2 (simple and adjusted) is provided for each regression. Heteroskedasticity-robust t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5% and 1% level respectively.

	(1)	(2)	(3)
	Dependent Variable: Flow _t		
Flow _{t-1}	0.72*** (12.60)	0.68*** (12.19)	0.69*** (12.21)
Return _t	0.02*** (7.47)		0.01 (0.92)
$\Delta(d - p)_t$		-2.47*** (-7.46)	-1.89** (-2.58)
Constant	0.02 (0.79)	0.09*** (3.14)	0.08** (2.12)
R^2	72.2	73.9	74.1
Adj. R^2	71.6	73.4	73.3

price ratio: $\Delta(d - p)_t$. If fund investors react to these news, then fund flows should be related to changes in dividend-price ratio.

This relationship is tested by regressing fund flows on its lag and concurrent changes in the dividend-price ratio. Results of this regression can be found in Table 3.3. Consistent with the information-response hypothesis, I find that an increase in dividend yield is linked to outflows from mutual funds. Moreover, the explanatory power of the dividend-price ratio is even higher than that of the market return: the adjusted R^2 of model (2) is 73.4 percent compared to the 71.6 percent of the baseline model (1). Including both market return and dividend-price ratio as regressors, as in column (3), leads to an insignificant coefficient for market return and no increase in adjusted R^2 . This result suggests that when explaining mutual fund flows the information of the variable market return is redundant when $\Delta(d - p)_t$ is included.

At first sight, the results of Table 3.3 do not seem surprising. Market returns and $\Delta(d - p)_t$ show a strong negative correlation (compare Table 3.1), since both variables

are to a large extent driven by price innovations. This negative correlation suggests that market returns and $\Delta(d - p)_t$ have the opposite effect on flows. But mutual fund flows are associated with changes in dividend-price ratio to an even greater extent than to market returns. If, as hypothesized, mutual fund flows react to macroeconomic news and their reaction to $\Delta(d - p)_t$ is stronger than to returns, then changes in dividend yield should also contain more information about the future economy. This is an additional hypothesis which can be tested.

Dividend yield is indeed a better forecasting variable for economic activity than market return. Just by looking at the correlation matrix of Table 3.1, we can see that the correlation between future GDP growth and $\Delta(d - p)_t$ is -0.42 , while the correlation between future GDP growth and market returns is only 0.32 . More specifically, in a forecasting regression for GDP growth, dividend yield achieves a higher adjusted R^2 than market return. Furthermore, in a joint forecasting model $\Delta(d - p)_t$ drives out market return, indicating that dividend yield contains more information about GDP growth and makes the information in returns redundant (see Appendix, Table B.1). The fact that dividend yield contains more macroeconomic information than market returns is consistent with the information-response hypothesis and explains why fund flows have a higher correlation with dividend yield than with returns.

3.5.2 Other predictive variables

The test of the information-response hypothesis is not restricted to the dividend-price ratio alone, but extends to further testable relations. If news about the real economy is the driving force behind mutual fund flow innovations, other variables that indicate riskier or less risky times should also be related to mutual fund flow innovations. Several other variables besides the dividend-price ratio relate to the equity premium, of which I investigate the following: default spread, term spread, relative T-Bill rate and the consumption-wealth ratio.

Default spread, term spread and the consumption-wealth ratio have been found to

Table 3.4:**Testable hypotheses: Predictive variables and mutual fund flows**

The table summarizes the findings for several predictive variables X_t and their connection to the equity premium, as well as their link to economic activity. It also displays the relation of a change in the predictive variable ΔX_t to mutual fund flows implied by the information-response hypothesis.

Variable:	Relation to economic activity:	Relation to equity premium:	Implied relation to mutual fund flows:
Dividend-price ratio	(-)	(+) ^a	(-)
Default spread	(-)	(+) ^b	(-)
Term spread	(-)	(+) ^c	(-)
Rel. T-Bill rate	(+)	(-) ^d	(+)
Consumption-wealth ratio	(-)	(+) ^e	(-)

^a Shiller, Fischer and Friedman (1984), Fama and French (1988), Campbell and Shiller (1988), and Ferson and Harvey (1991)

^b Fama and French (1989), and Chen (1991)

^c Campbell (1987), Fama and French (1989), and Chen (1991)

^d Campbell (1991) and Hodrick (1992)

^e Lettau and Ludvigson (2001, 2005)

be *positively* related to the equity premium (see Fama and French 1988, Campbell and Shiller 1988, Fama and French 1989, Chen 1991, Lettau and Ludvigson 2001), while the relative T-Bill rate has been found to be *negatively* related to the equity premium (see Campbell 1991, Hodrick 1992). Consequently, under the information-response hypothesis an increase in default and term spread as well as *cay* should be associated with *outflows*, and an increase in the relative T-Bill rate should be accompanied by *inflows* from mutual fund investors. As discussed before, it is the unexpected *change* in predictive variables, which reflects news, and if fund investors respond to this information then we expect a co-movement of the first difference of predictive variables and fund flows.

Table 3.4 summarizes the literature on return predictability and the testable hypotheses for mutual fund flows under the information-response hypothesis. It provides the predictive variables mentioned above, their link to the business cycle and their relation to expected

returns. The last column shows the testable relation of predictive variables to mutual fund flows implied by the information-response hypothesis assuming that mutual fund investors are less willing to hold equity in poor economic times. Under the information-response hypothesis the positive correlation of flows and market returns implies that they sell at bad news and buy at good news, i.e. that they are less willing to hold equity in bad times.⁸

Isolated from the macroeconomic context, one might wonder why variables that signal high expected returns should result in outflows from equity funds. Should mutual fund investors not react to the signal of high expected returns and buy equity? Under the information-response hypothesis it is exactly the other way around. It is news about riskier economic times that is reflected in the predictive variables. As a response to this news, mutual fund investors (possibly along with other investors) reduce their equity holdings. And other investors who are willing to hold equity in riskier economic times are compensated by higher expected returns.

As before the regression analysis results of flows on a set of contemporaneous predictive variables should not be interpreted causally. We also do not assume that fund investors observe and respond to these variables.⁹ Rather we see the variables as proxies for news, because they reflect expectations about the future. Thus, the following analysis investigates how fund flows are related to these news proxies and whether fund flows are better described by these news variables than simply by market returns.

The regression results of mutual fund flows on other predictive variables are presented in Table 3.5. Panel A displays the predictive variables without the change in dividend yield, Panel B the predictive variables combined with the change in dividend yield. The results in Panel A show that mutual fund flow innovations are negatively related to changes in default spread. An increase in default risk signals riskier times to invest and thus is associated with outflows from mutual funds. The opposite is the case for the relative T-Bill rate, where a rise in the relative T-Bill rate indicates a lower equity premium: more

⁸See footnote 2 for details.

⁹ While in principle it would be possible for an investor to observe dividend-price ratio, default spread, term spread and T-Bill rate in real time, the consumption-wealth ratio cannot be observed in real time due to a delayed release of macroeconomic variables.

investors are willing to hold equity, which results in higher inflows. Mutual fund flow innovations seem to be unrelated to changes in term spread. This can be explained by the fact that term spread is related to a greater degree to past and contemporaneous economic activity than to future economic activity (see, e.g., Fama and French 1989, Fama 1991, or Panel B of Table 3.1). Thus, term spread is rather an indicator of bad times than a proxy for news about imminent bad times, explaining why fund flows show no relation to it.

Finally, mutual fund flows are, as predicted, negatively linked to the consumption-wealth ratio. The consumption-wealth ratio is high before and around economic contractions and therefore positively related to the equity premium. Increases in *cay* signal poor economic times, which are accompanied by a downward adjustment in mutual fund investors' equity holdings. Overall, these findings support the information-response hypothesis: bad news about the economy (reflected in a rise in default spread and consumption-wealth ratio) leads to outflows by mutual fund investors, while good news about the economy (indicated by an increase in relative T-Bill rate) leads to inflows.

Panel B in Table 3.5 constitutes an investigation of which predictive variables have an influence on mutual fund flows in addition to the dividend-price ratio. The results in column (3) show that the relative T-Bill rate provides additional explanation for mutual fund flows with an adjusted R^2 of 74.5, which is higher than the model including only dividend yield (Adj. R^2 : 73.4, see Table 3.3). The default spread, on the other hand, becomes insignificant when the dividend-price ratio is added. This is not surprising, since it is well established that default spread has no marginal explanatory power for expected returns, when the dividend-price ratio is included (see, e.g., Fama and French 1989, Chen 1991, Hodrick 1992). This is due to the fact that the two variables contain similar information about the business cycle. (See Table 3.1, Panel B for the correlation structure of these variables.) If a variable has no additional information with respect to the equity premium, then it should not have an additional effect on mutual fund flows either.

The consumption-wealth ratio *cay*, on the other hand, is known to provide another independent dimension to the predictability of excess returns (Lettau and Ludvigson 2001,

Table 3.5:
Mutual fund flows and changes in other predictive variables

This table shows the regression results of mutual fund flows on lagged flows and changes in predictive variables, where the expected coefficient signs are given in parentheses (compare Table 3.4). $\Delta(d - p)_t$ is the change in log dividend-price ratio, $\Delta\text{Default}_t$ the change in default spread, ΔTerm_t the change in term spread, $\Delta\text{Rel. T-Bill}_t$ the change in the relative 3-month T-Bill rate, and Δcay_t the change in the consumption-wealth ratio. The table provides R^2 and adjusted R^2 for each regression. Heteroskedasticity-robust t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5% and 1% level respectively.

	(1)	(2)	(3)	(4)	(5)
Panel A:					
Dependent Variable: Flow_t					
Flow_{t-1}	0.72*** (12.43)	0.72*** (10.48)	0.73*** (11.22)	0.76*** (13.33)	0.76*** (14.09)
$\Delta(d - p)_t$ (-)					
$\Delta\text{Default}_t$ (-)	-0.27** (-2.11)				-0.07 (-0.69)
ΔTerm_t (-)		-0.05 (-0.56)			0.06 (0.79)
$\Delta\text{Rel. T-Bill}_t$ (+)			0.18*** (3.47)		0.15** (2.53)
Δcay_t (-)				-21.69*** (-5.78)	-19.80*** (-4.94)
Return_t					
Constant	0.10*** (2.72)	0.10** (2.35)	0.09** (2.37)	0.08*** (2.80)	0.08*** (2.84)
R^2	55.1	52.8	56.9	67.5	70.4
Adj. R^2	54.2	51.8	56.0	66.9	68.9

Table 3.5 -Continued

		(1)	(2)	(3)	(4)	(5)	(6)
Panel B:							
Dependent Variable: Flow _t							
Flow _{t-1}		0.69*** (11.63)	0.68*** (11.82)	0.69*** (13.10)	0.71*** (13.49)	0.72*** (14.04)	0.73*** (14.01)
$\Delta(d-p)_t$	(-)	-2.59*** (-7.88)	-2.46*** (-7.36)	-2.34*** (-7.13)	-1.91*** (-4.91)	-1.93*** (-5.34)	-1.67** (-2.37)
$\Delta\text{Default}_t$	(-)	0.09 (0.66)				0.12 (1.08)	0.11 (0.97)
ΔTerm_t	(-)		-0.02 (-0.40)			0.07 (1.06)	0.07 (1.08)
$\Delta\text{Rel. T-Bill}_t$	(+)			0.10*** (2.95)		0.12*** (2.98)	0.13*** (2.72)
Δcay_t	(-)				-10.66** (-2.61)	-10.58** (-2.51)	-10.15** (-2.26)
Return _t							0.00 (0.39)
Constant		0.09*** (2.98)	0.10*** (3.14)	0.09*** (3.30)	0.09*** (3.30)	0.08*** (3.24)	0.07** (2.47)
R^2		74.1	74.0	75.3	76.5	78.2	78.2
Adj. R^2		73.4	73.2	74.5	75.7	76.8	76.6

2005). If *cay* contains additional information about the risk premium of the stock market, changes in *cay* should also help to explain unexpected fund flows in addition to the dividend-price ratio. The results documented in column (4) suggest that this is the case. The adjusted R^2 of this model is 75.7, which is considerably higher than that of the benchmark model using only lagged flows and contemporaneous returns as explanatory variables, which has an adjusted R^2 of 71.6 percent (see Table 3.2, column (3)). The joint model uniting all predictive variables even yields an adjusted R^2 of 76.8 percent, as can be seen in column (5). These results are also robust, when market return is included as an additional explanatory variable as in column (6). The predictive variables stay significant, while market return adds no explanatory power to fund flows. These results are even more pronounced, if we only consider unexpected flows as done in Table 3.6. While market return only explains 40.8 percent of the variation of unexpected flows, a model of all predictive variables yields an adjusted R^2 of 51.7 percent.

Evidence provided in Table 3.5 and 3.6 shows that mutual fund flows are strongly related to economic fundamentals, which stands in contrast to the feedback-trader and price-pressure hypothesis. The key insight is: mutual fund flows are *better* described by macroeconomic news proxies than simply by market returns. This means: they are not merely a feedback response to market returns; they are not merely uninformed investors that induce price pressure.

Figure 3.1 depicts the relation between mutual fund flows and the most important predictive variables: dividend-price and the consumption-wealth ratio. As mentioned before, rises in the dividend-price ratio as well as *cay* occur at the beginning of and during recessions, coinciding with outflows from equity funds. The figure also clarifies why *cay* provides additional explanatory power with regard to mutual fund flows. On a number of occasions where we observe outflows, such as the recessions of 1990/1991 and 2001, *cay* is more responsive than the dividend-price ratio, thus providing additional information.

Table 3.6:
Unexpected fund flows and changes in predictive variables

This table shows the regression results of unexpected mutual fund flows on lagged flows and changes in predictive variables. For details on the variables see Table 3.5. The table provides R^2 and adjusted R^2 for each regression. Heteroskedasticity-robust t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5% and 1% level respectively.

		(1)	(2)	(3)
		Dependent Variable: Unexpected Flow _t		
$\Delta(d - p)_t$	(-)	-1.93***	-1.66**	
		(-5.20)	(-2.34)	
$\Delta\text{Default}_t$	(-)	0.12	0.11	
		(1.10)	(0.98)	
ΔTerm_t	(-)	0.07	0.07	
		(1.10)	(1.11)	
$\Delta\text{Rel. T-Bill}_t$	(+)	0.12***	0.13***	
		(3.02)	(2.79)	
Δcay_t	(-)	-10.63**	-10.16**	
		(-2.53)	(-2.27)	
Return _t		0.02***		0.00
		(8.45)		(0.40)
Constant		-0.07***	-0.01	-0.02
		(-2.95)	(-0.57)	(-0.76)
R^2		41.4	54.1	54.1
Adj. R^2		40.8	51.7	51.3

3.6 Mutual fund flows and future economic activity

3.6.1 Vector autoregression analysis

The previous section investigated the relation of predictive variables to mutual fund flows, while stressing the link of both to the real economy. Now I will analyze in detail the relation of mutual fund flows to the real economy, as an alternative test to see whether mutual fund flows react to macroeconomic news. The idea behind this test is that if mutual fund flows respond to news about the real economy, then mutual fund flows should be able to predict real economic activity. If there is news about a worsening economy, the marginal mutual fund investor, unwilling to hold equity funds through this time, will withdraw his or her shares. On the other hand, if positive news about the economy occurs, the marginal investor will be more willing to hold equity funds, increasing his or her shares. If mutual fund investors are *on average* right, then the state of the economy should be worse after outflows and better after inflows into mutual funds (compare Roll 1984). This is the second main hypothesis implied by the information-response explanation of mutual fund flows.

This hypothesis will first be tested within a bivariate vector autoregression (VAR) framework of mutual fund flows and economic activity growth. Measures for economic activity are real GDP, industrial production, consumption and labor income. To answer the question of whether mutual fund flows contain information about future economic activity, I employ the concept of Granger causality. That is, I test whether lags of economic activity provide statistically significant information about future mutual fund flows, or whether lags of mutual fund flows provide statistically significant information about future economic activity. Of course, Granger causality does not imply true causality, i.e., it does not say that mutual fund flows cause higher economic activity or vice versa. It merely states that one variable contains information about the other. And for that matter it is exactly the question we are interested in, since we want to investigate if mutual fund flows react to macroeconomic news and therefore contain information about the real economy.

Table 3.7 shows the estimation results of the VAR model using one lag. The small

Table 3.7:
Mutual fund flows and economic activity

The table provides estimates of a VAR(1) of mutual fund flows and proxies for economic activity growth. Measures for economic activity are real gross domestic product (Panel A), industrial production (Panel B), consumption (Panel C) and labor income (Panel D). It also displays a Granger causality test for fund flows and economic activity. In column (1) the Granger causality F-statistic tests that flows are excluded from the economic activity growth equation, and in column (2) that economic activity growth is excluded from the flow equation. Heteroskedasticity-robust t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5% and 1% level respectively.

	(1)	(2)
	Dependent Variables:	
Panel A:	GDP	
	Growth _t	Flow _t
GDP Growth _{t-1}	0.35*** (3.94)	-0.02 (-0.29)
Flow _{t-1}	0.43*** (3.42)	0.74*** (10.00)
Constant	0.29*** (3.73)	0.10** (2.18)
Granger Causality: F-Statistic	11.7	0.1
p-value	0.00	0.77
	Dependent Variables:	
Panel B:	Industrial Production	
	Growth _t	Flow _t
Ind. Production Growth _{t-1}	0.35*** (4.07)	0.02 (0.82)
Flow _{t-1}	1.17*** (4.17)	0.70*** (9.35)
Constant	-0.11 (-0.77)	0.09** (2.44)
Granger Causality: F-Statistic	17.3	0.7
p-value	0.00	0.41

Table 3.7 -Continued

	(1)	(2)
	Dependent Variables:	
Panel C:	Consumption	
	Growth _t	Flow _t
Consumption Growth _{t-1}	0.29*** (3.10)	0.03 (0.50)
Flow _{t-1}	0.30** (2.50)	0.72*** (9.92)
Constant	0.43*** (4.98)	0.08 (1.45)
Granger Causality: F-Statistic	6.2	0.2
p-value	0.01	0.62
	Dependent Variables:	
Panel D:	Labor Income	
	Growth _t	Flow _t
Labor Income Growth _{t-1}	0.36*** (4.02)	0.03 (0.88)
Flow _{t-1}	0.48*** (2.86)	0.71*** (9.70)
Constant	0.20** (2.09)	0.08* (1.97)
Granger Causality: F-Statistic	8.2	0.8
p-value	0.01	0.38

lag length is chosen in order to provide a parsimonious model of mutual fund flows and economic activity growth. This model selection is also supported by the Schwarz-Bayes information (SBIC) criterion. For VAR models including additional lags, see Table B.2 in the Appendix.

For all four proxies of economic activity we find a consistent pattern: mutual fund flows help to predict economic activity growth, but economic activity growth does not help to predict mutual fund flows. In the economic activity equation, lagged flows are significant for all proxies of economic activity, while in the fund flow equation lagged economic activity is insignificant. This result is supported by the Granger causality F-test. The Granger causality test results are robust for VAR models of several lag lengths (see Appendix, Table B.2).¹⁰ This result provides further evidence for the information-response hypothesis. Mutual fund flows are not unrelated to fundamentals but rather react to macroeconomic news and therefore possess predictive power for the real economy.

Table 3.7 documents a new and important fact about mutual fund flows: they are forward-looking. At a first glance, this might seem surprising, yet we are quite familiar with the fact that financial variables are forward-looking: it is well established, for example by Fama (1990) and Schwert (1990), that stock prices or returns predict economic activity. If fund flows and returns react to the same macroeconomic information, it must follow that mutual fund flows predict economic activity as well. Hence, the next section considers the joint forecasting ability of market returns and fund flows.

3.6.2 Forecasting comparison of market returns and fund flows

Following Ludvigson (2004) and Lemmon and Portniaguina (2006) I run a forecasting regression of economic activity on its lags and lagged market return and/or flows. The baseline model is a simple model of economic activity growth regressed on its four lags, the results of which are not reported for reasons of brevity. I calculate the increment of

¹⁰Vector autoregression models are estimated with lags 1 through 4. For all VAR models, Granger causality tests yield similar results.

adjusted R^2 , which is the percentage point increase of adjusted R^2 relative to the baseline model.

Table 3.8 shows the results of this forecasting regression for real GDP, industrial production, consumption and labor income growth. As documented in previous literature (e.g. Fama 1990, Schwert 1990), we see that stock market returns help to predict economic activity in addition to its lagged values. The incremental R^2 varies between 3.5 and 10.2 percent depending on the economic activity measure considered, and mutual fund flows predict economic activity in a similar manner. The regression coefficient is significantly different from zero throughout all specifications and the incremental R^2 is comparable to that of the market return. Consumption and labor income growth are in general harder to predict. Their regressions' incremental R^2 is lower for both market return and fund flows.

When both market return and mutual fund flows are included to predict economic activity, we observe a reduction in regression coefficients and significance for both variables indicating that returns and flows contain partly redundant information about future economic activity. This is especially the case for GDP and consumption growth, where the market return coefficient becomes insignificant. Overall, these results imply that market returns and mutual fund flows contain similar (but not completely identical) information about the real economy, which explains their co-movement over time.

Table 3.8:
Mutual fund flows, market returns, and real economic activity

The table provides the estimates of a forecasting regression of economic activity growth. Measures for economic activity are gross domestic product (GDP), industrial production, consumption and labor income. The forecasting regression also includes four lags of the dependent variable (baseline model). The incremental adjusted R^2 (in percent) is reported relative to the baseline model, which includes only lagged values of the dependent variable. Heteroskedasticity-robust t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5% and 1% level respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent Variable: GDP Growth _t			Dependent Variable: Ind. Production Growth _t		
Return _{t-1}	0.02*** (2.80)		0.01 (1.63)	0.06*** (3.78)		0.04*** (2.70)
Flow _{t-1}		0.40*** (3.30)	0.31** (2.33)		1.14*** (4.50)	0.77*** (3.09)
Incremental Adj. R^2	4.4	6.4	7.4	10.2	9.1	13.2
	Dependent Variable: Consumption Growth _t			Dependent Variable: Labor Income Growth _t		
Return _{t-1}	0.01* (1.81)		0.01 (1.42)	0.02** (2.45)		0.02* (1.80)
Flow _{t-1}		0.24** (2.13)	0.15 (1.55)		0.49*** (2.82)	0.34** (2.24)
Incremental Adj. R^2	3.5	2.6	3.8	5.4	5.3	7.0

3.7 Concluding remarks

The aim of this paper was to test competing theories about the co-movement of mutual fund flows and stock market returns: feedback-trading, price-pressure or common response to information. Results presented in this paper provide evidence for the theory that stock market returns and mutual fund investors commonly react to macroeconomic information. Mutual fund flows are better described by variables that proxy for macroeconomic information than by stock market return alone. In particular, flows into equity funds are related to changes in dividend-price ratio and the consumption-wealth ratio. The information response hypothesis is further supported by the fact that mutual fund flows in themselves contain information: mutual fund flows - like stock returns - are forward-looking and help to predict real economic activity.

These results raise interesting questions about market timing, tactical asset allocation and return predictability. Moreover, by investigating the portfolio choices of one particular investor group, mutual fund investors, this analysis is a first step away from the average investor view towards looking at different investor preferences as proposed by Cochrane (2011). While an analysis of other investor groups would be beyond the scope of this paper, it is an interesting avenue for future research.

Appendix B: Additional Tables

Table B.1:
Economic activity forecasting comparison:
Market return and change in dividend yield

The table provides the estimates of a forecasting regression of economic activity proxied by GDP growth (see Table 3.8). The forecasting regression also includes four lags of the dependent variable (baseline model). The incremental adjusted R^2 (in percent) is reported relative to the baseline model, which includes only lagged values of the dependent variable. Heteroskedasticity-robust t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5% and 1% level respectively.

	(1)	(2)	(3)
	GDP Growth _t		
Return _{t-1}	0.02*** (2.80)		-0.02 (-1.06)
(d - p) _{t-1}		-2.23*** (-4.17)	-3.90** (-2.47)
Incremental Adj. R^2	4.4	7.5	7.6

Table B.2: Mutual fund flows and real economic activity: further specifications

In this table the analysis of Table 3.7 is repeated for further lag lengths: VAR(1) - VAR(4). The table further provides the Schwarz-Bayes (SBIC) and Akaike (AIC) information criterion for model selection (minimum in bold face). For each equation the table also provides the Ljung-Box Q-statistic for the test that residuals are not autocorrelated (up to lag 20). Heteroskedasticity-robust t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5% and 1% level respectively.

Panel A:	VAR(1)		VAR(2)		VAR(3)		VAR(4)	
	GDP Growth _t	Flow _t	GDP Growth _t	Flow _t	GDP Growth _t	Flow _t	GDP Growth _t	Flow _t
GDP Growth _{t-1}	0.35*** (3.94)	-0.02 (-0.29)	0.23*** (2.38)	-0.02 (-0.32)	0.28*** (2.72)	0.01 (0.12)	0.27*** (2.66)	0.02 (0.26)
GDP Growth _{t-2}			0.29*** (3.22)	-0.01 (-0.24)	0.37*** (3.87)	-0.04 (-0.71)	0.35*** (3.28)	-0.04 (-0.63)
GDP Growth _{t-3}					-0.16* (-1.68)	-0.04 (-0.62)	-0.15 (-1.45)	-0.06 (-0.92)
GDP Growth _{t-4}							0.06 (0.64)	0.04 (0.68)
Flow _{t-1}	0.43*** (3.42)	0.74*** (10.00)	0.60*** (3.74)	0.68*** (6.84)	0.59*** (3.72)	0.66*** (6.63)	0.62*** (3.82)	0.64*** (6.39)
Flow _{t-2}			-0.28* (-1.66)	0.08 (0.79)	-0.20 (-1.03)	-0.05 (-0.44)	-0.20 (-1.02)	-0.05 (-0.42)
Flow _{t-3}					-0.13 (-0.77)	0.23*** (2.15)	-0.05 (-0.27)	0.14 (1.08)
Flow _{t-4}							-0.15 (-0.80)	0.12 (1.11)
Constant	0.29*** (3.73)	0.10** (2.18)	0.21** (2.58)	0.10** (2.05)	0.25*** (3.05)	0.10** (2.03)	0.24*** (2.63)	0.08 (1.45)
SBIC		2.20		2.27		2.37		2.54
AIC		2.05		2.01		2.01		2.07
Q-Statistic	24.2	19.9	19.4	21.5	21.9	18.2	20.0	16.8
p-value	0.23	0.46	0.49	0.37	0.34	0.58	0.46	0.67
Granger Causality: F-Statistic	11.7	0.1	7.6	0.1	4.9	0.5	3.9	0.4
p-value	0.00	100.00	0.00	0.88	0.00	0.71	0.01	0.82

Table B.2 -Continued

Panel B:	VAR(1)		VAR(2)		VAR(3)		VAR(4)	
	Industrial Prod.	Flow _t	Industrial Prod.	Flow _t	Industrial Prod.	Flow _t	Industrial Prod.	Flow _t
	Growth _t	Flow _t	Growth _t	Flow _t	Growth _t	Flow _t	Growth _t	Flow _t
Ind. Production Growth _{t-1}	0.35*** (4.07)	0.02 (0.82)	0.35*** (3.58)	0.02 (0.85)	0.35*** (3.54)	0.03 (0.98)	0.36*** (3.57)	0.03 (1.05)
Ind. Production Growth _{t-2}			0.00 (0.03)	-0.01 (-0.53)	0.00 (0.03)	-0.02 (-0.69)	0.02 (0.15)	-0.02 (-0.68)
Ind. Production Growth _{t-3}					0.05 (0.45)	-0.03 (-1.05)	0.09 (0.85)	-0.03 (-1.09)
Ind. Production Growth _{t-4}							-0.11 (-1.06)	-0.00 (-0.06)
Flow _{t-1}	1.17*** (4.17)	0.70*** (9.35)	1.29*** (3.46)	0.68*** (6.79)	1.31*** (3.48)	0.66*** (6.71)	1.28*** (3.31)	0.63*** (6.34)
Flow _{t-2}			-0.18 (-0.46)	0.05 (0.48)	-0.09 (-0.19)	-0.07 (-0.60)	-0.13 (-0.28)	-0.07 (-0.54)
Flow _{t-3}					-0.22 (-0.55)	0.24** (2.30)	-0.07 (-0.14)	0.13 (1.05)
Flow _{t-4}							-0.09 (-0.21)	0.16 (1.54)
Constant	-0.11 (-0.77)	0.09** (2.44)	-0.10 (-0.65)	0.09** (2.25)	-0.08 (-0.49)	0.07* (1.73)	-0.04 (-0.25)	0.06 (1.47)
SBIC		3.77		3.96		4.11		4.27
AIC		3.62		3.71		3.74		3.80
Q-Statistic	11.8	17.8	12.3	18.6	12.4	15.1	9.7	14.4
p-value	0.92	0.60	0.90	0.55	0.90	0.77	0.97	0.81
Granger Causality: F-Statistic	17.3	0.7	8.4	0.4	5.5	1.0	3.8	0.8
p-value	0.00	0.41	0.00	0.68	0.00	0.41	0.01	0.53

Table B.2 - Continued

Panel C:	VAR(1)		VAR(2)		VAR(3)		VAR(4)	
	Cons. Growth _t	Flow _t	Cons. Growth _t	Flow _t	Cons. Growth _t	Flow _t	Cons. Growth _t	Flow _t
Consumption Growth _{t-1}	0.29*** (3.10)	0.03 (0.50)	0.20** (2.09)	0.03 (0.42)	0.12 (1.26)	0.01 (0.16)	0.14 (1.39)	0.06 (0.89)
Consumption Growth _{t-2}			0.26*** (2.76)	-0.01 (-0.09)	0.19** (2.12)	-0.02 (-0.28)	0.19** (2.03)	0.00 (0.03)
Consumption Growth _{t-3}					0.38*** (4.14)	0.00 (0.06)	0.39*** (4.18)	0.01 (0.11)
Consumption Growth _{t-4}							-0.06 (-0.62)	-0.10 (-1.50)
Flow _{t-1}	0.30** (2.50)	0.72*** (9.92)	0.18 (1.14)	0.67*** (6.74)	0.23 (1.57)	0.67*** (6.70)	0.22 (1.48)	0.64*** (6.41)
Flow _{t-2}			0.11 (0.67)	0.06 (0.61)	0.24 (1.37)	-0.06 (-0.48)	0.24 (1.36)	-0.07 (-0.56)
Flow _{t-3}					-0.31** (-2.09)	0.19* (1.82)	-0.31* (-1.67)	0.09 (0.71)
Flow _{t-4}							0.01 (0.09)	0.16 (1.51)
Constant	0.43*** (4.98)	0.08 (1.45)	0.30*** (3.23)	0.08 (1.28)	0.16* (1.78)	0.07 (1.15)	0.18* (1.88)	0.09 (1.42)
SBIC		2.13		2.22		2.21		2.34
AIC		1.97		1.96		1.84		1.87
Q-Statistic	29.9	20.0	25.4	21.1	9.5	17.1	8.3	14.9
p-value	0.07	0.46	0.18	0.39	0.98	0.65	0.99	0.78
Granger Causality: F-Statistic	6.2	0.2	2.6	0.1	3.2	0.0	2.3	0.6
p-value	0.01	0.62	0.08	0.91	0.03	0.99	0.07	0.66

Table B.2 -Continued

	VAR(1)		VAR(2)		VAR(3)		VAR(4)	
	Labor Income Growth _t	Flow _t	Labor Income Growth _t	Flow _t	Labor Income Growth _t	Flow _t	Labor Income Growth _t	Flow _t
Labor Income Growth _{t-1}	0.36*** (4.02)	0.03 (0.88)	0.25*** (2.57)	0.04 (0.92)	0.19* (1.89)	0.04 (0.94)	0.13 (1.32)	0.06 (1.33)
Labor Income Growth _{t-2}			0.23*** (2.43)	-0.02 (-0.47)	0.17* (1.68)	-0.02 (-0.40)	0.12 (1.25)	-0.01 (-0.17)
Labor Income Growth _{t-3}					0.19* (1.97)	-0.06 (-1.37)	0.13 (1.31)	-0.05 (-1.11)
Labor Income Growth _{t-4}							0.25*** (2.34)	-0.07 (-1.57)
Flow _{t-1}	0.48*** (2.86)	0.71*** (9.70)	0.25 (1.13)	0.67*** (6.73)	0.29 (1.30)	0.65*** (6.66)	0.36 (1.62)	0.60*** (6.10)
Flow _{t-2}			0.26 (1.16)	0.06 (0.57)	0.17 (0.66)	-0.05 (-0.44)	0.21 (0.81)	-0.05 (-0.43)
Flow _{t-3}					0.07 (0.32)	0.21** (2.04)	-0.12 (-0.43)	0.13 (1.09)
Flow _{t-4}							0.14 (0.60)	0.15 (1.48)
Constant	0.20** (2.09)	0.08* (1.97)	0.11 (1.11)	0.08* (1.92)	0.05 (0.53)	0.09** (2.02)	-0.04 (-0.39)	0.11** (2.35)
SBIC		2.80		2.91		3.01		3.09
AIC		2.65		2.65		2.65		2.62
Q-Statistic	17.8	18.3	15.5	18.3	13.4	15.4	7.7	17.4
p-value	0.60	0.57	0.75	0.56	0.86	0.75	0.99	0.62
Granger Causality: F-Statistic	8.2	0.8	4.0	0.4	2.7	0.9	2.6	1.4
p-value	0.01	0.38	0.02	0.65	0.05	0.45	0.04	0.24

Chapter 4

Can internet search queries help to predict stock market volatility?*

Abstract

This paper studies the dynamics of stock market volatility and retail investor attention measured by internet search queries. We find a strong co-movement of stock market indices' realized volatility and the search queries for their names. Furthermore, Granger causality is bi-directional: high searches follow high volatility, and high volatility follows high searches. Using the latter feedback effect to predict volatility we find that search queries contain additional information about market volatility. They help to improve volatility forecasts in-sample and out-of-sample as well as for different forecasting horizons. Search queries are particularly useful to predict volatility in high-volatility phases.

*This chapter is based on the working paper "*Can internet search queries help to predict stock market volatility?*" by Dimpfl, T. and S. Jank (2011).

4.1 Introduction

Large stock market movements capture investors' attention. This can be seen in Figure 4.1, which depicts a strong co-movement between volatility of four leading stock market indices (Dow Jones, FTSE, CAC and DAX) and Google search queries for their name in their home country. For example, when volatility of the Dow Jones spiked at an almost record high of over 150% annualized on October 10, 2008, the number of submitted searches for Dow Jones rose to more than eleven times the average.

Internet search queries can be interpreted as a measure for retail investors' attention to the stock market as recently suggested by Da, Engelberg and Gao (2011). While professional investors monitor the leading index all the time, retail investors are likely not to do so. Once the latter perceive an increased demand for information about the stock index, they are likely to use the internet as a source of information.

In this paper we study in detail the dynamics of retail investor attention for the aggregate stock market, proxied by internet searches, and stock market volatility. The key finding of this paper is that there exists bi-directional Granger causality between realized volatility of the stock market indices Dow Jones, FTSE, CAC and DAX and search activity for their respective names. Most importantly, search query data have predictive power for future volatility of the stock market. We exploit this finding and augment various models of realized volatility with search query data. The forecasting precision can be significantly improved when data on search queries enter the prediction equation. The improvement is evident both for in-sample as well as for out-of-sample forecasts. The longer the forecast horizon, the more efficiency gains are apparent. Furthermore, the data on internet search queries help to predict volatility more accurately in periods of high volatility, i.e. when a precise prediction is vital.

These findings contribute to our knowledge of stock market volatility and its long memory characteristics documented for example by Andersen and Bollerslev (1997). In particular, the findings are consistent with agent-based models of stock market volatility (e.g. Lux and Marchesi 1999, Alfarano and Lux 2007). In the model by Lux and Marchesi

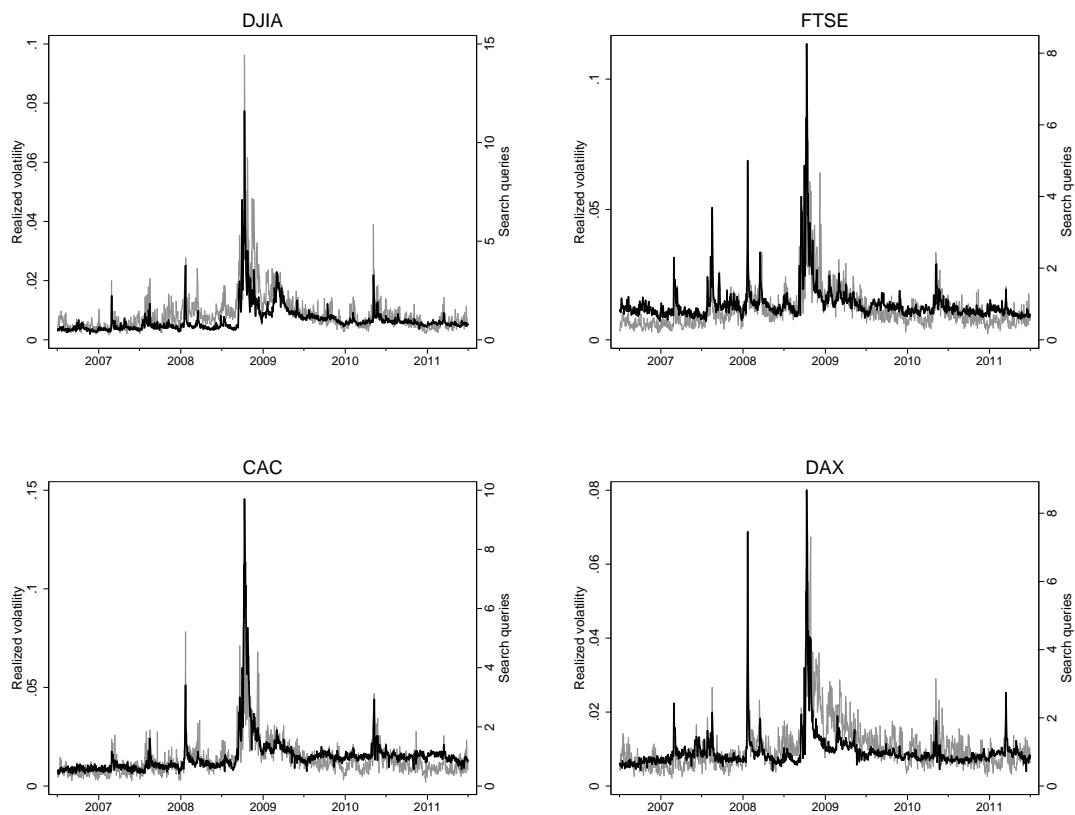


Figure 4.1:
Realized volatility and search activity

This figure displays daily realized volatilities (gray) and search queries (black) of the stock indices DJIA, FTSE, CAC and DAX from July 1, 2006 to June 30, 2011. Search queries are standardized, such that the sample average equals one.

(1999) noise traders are seen as a source of additional volatility in the stock market. A fundamental shock in volatility triggers noise trading, which in turn causes volatility. Taking internet searches as a measure of retail investors' attention, we observe exactly this pattern of high volatility followed by high retail investor attention, which is then followed by high volatility. Our results are also in line with recent empirical evidence by Foucault et al. (2011), who - drawing on a natural experiment in France - find that retail investors' trading activity leads to a higher level of volatility in individual stocks.

A natural question which arises is how much of a stock market's volatility is driven by noise traders and how much is fundamental. In a long-run variance decomposition we

find that log search queries account for 9% to 23% of the variance of log stock market volatility.¹ However, this share has to be interpreted with caution. Although, internet search queries are most likely a proxy for retail investors' attention we do not observe whether the individuals searching for the index are the same that actually trade and cause the higher volatility. Still, irrespective of the link between search queries and noise traders, the fact that retail investor attention contains information about future volatility can be used to improve volatility forecasts, which is the main focus of this paper.

In a forecasting context, other recent studies have successfully used Google search volume data. For example Ginsberg et al. (2009) use search query data to predict influenza epidemics and Choi and Varian (2009a) and Choi and Varian (2009b) employ Google search data to forecast unemployment rates and retail sales, respectively. In the field of finance search query data are used to measure retail investor attention (Bank, Larch and Peter 2011, Da et al. 2011, Jacobs and Weber forthcoming) and to predict earnings (Da, Engelberg and Gao 2010a, Drake, Roulstone and Thornock 2011). Da, Engelberg and Gao (2010b) use search queries related to household concerns to measure investor sentiment.

We proceed as follows. In Section 4.2 we describe our data set of realized volatilities and search engine data. Section 4.3 presents standard models for predicting volatility and highlights the contribution of search query data in the modeling process. Section 4.4 evaluates in- and out-of-sample forecasts of realized volatility and Section 4.5 concludes.

4.2 Data and descriptive statistics

Our analysis focuses on the US stock market index and three major European indices from July 2006 to June 2011: the Dow Jones Industrial Average (DJIA), the FTSE 100, the CAC 40 and the DAX. European intraday market index prices are obtained from Tick Data while US intraday prices are provided by RC Research Price-Data.

We construct a time series of daily realized volatilities $RV_{i,t}$ as introduced by Andersen,

¹A similar share is found by Foucault et al. (2011) even though using a different sample period. They estimate that retail investors contribute to about 23% of the volatility in stock returns.

Bollerslev, Diebold and Labys (2003) for the four stock indices i the following way:

$$RV_{i,t} = \sqrt{\sum_{j=1}^n r_{i,t,j}^2}, \quad (4.1)$$

where $r_{i,t,j}^2$ are squared intraday log-price changes of index i on day t during interval j and n is the number of such intraday return intervals. We compute these price changes over 10 minute intervals in order to circumvent the well-documented microstructure effects (see e.g. Andersen et al. 2003, Andersen, Bollerslev and Meddahi 2011, Ghysels and Sinko 2011).²

Descriptive statistics of the realized volatilities are presented in the upper panel of Table 4.1. As is evident from the skewness and kurtosis measures, the volatility time series are heavily skewed and far from being normally distributed. We therefore resort to the log of the realized volatility as, amongst others, suggested by Andersen, Bollerslev, Diebold and Ebens (2001) and Andersen et al. (2003). The lower panel of Table 4.1 shows that, even though normality of the data still has to be rejected, the data are by far better behaved than before the transformation; in particular excess kurtosis is significantly reduced. Figure 4.2 holds the autocorrelation functions for realized volatilities of the indices DJIA, FTSE, CAC and DAX. The plots reveal the well-known pattern of only slowly decaying autocorrelations (compare e.g. Andersen et al. 2001).

The data on Google search queries are obtained through *Google Trends*.³ We use daily data on search volume from July 2006 to June 2011 for the keywords “Dow” (US search queries), “FTSE” (UK search queries), “CAC” (search queries in France) and “DAX” (search queries in Germany) within the respective countries. Before July 2006 search volume data at daily frequency exhibit many missing values. We therefore start our sample in the

²To exclude the possibility that our results are driven by the sampling frequency, we also compute realized volatility over 5 and 15 minute intervals. Our results are robust to this alteration.

³Source: <http://www.google.com/trends>.

second half of 2006.⁴ To match searches to the respective time series of realized volatility we only consider trading days of the stock markets in question.

An important issue when measuring the investors' attention for a certain index is that stock indices often go by many names. The question which search term individuals use when looking for information about the stock market is answered most easily for the UK, France and Germany, since the leading indices' names are only few. In general, the short name of the index is preferred. The number of search queries of "FTSE 100" amounts to approximately 45% of the searches for "FTSE", and queries for "CAC 40" to about 77% of queries for "CAC". The term "DAX 30" is less commonly used in Germany and search volumes are negligible. Correlations between the different search terms are high with 0.95 for "FTSE 100" and "FTSE", and 0.998 for "CAC" and "CAC 40".

In the US, the picture is similar even though the Dow Jones is known under a variety of names and acronyms. We find that the most widely used search term is simply "Dow", followed by "Dow Jones" which amounts to approximately 45% of the search volume of "Dow". Searches of the full name "Dow Jones Industrial Average" amount to 10% when compared to "Dow", search queries for ticker symbols such as "DJIA" and "DJI" to 17% and 7% respectively. Even though the magnitude of searches is quite different, the correlation between the search queries is remarkably high. The pairwise correlation of the named terms is in all cases above 0.97.⁵ Since the correlation between the various index names is consistently very high, we use the search term that is mostly used.

For the US we use the Dow Jones as leading index. An alternative index would be the S&P 500, which is commonly modeled in the realized volatility literature. However, the S&P 500 is less suited for our purposes, because it is less followed by retail investors. We find that the S&P 500 overall attracts less attention than the Dow Jones. In our sample period the search term "Dow" has been submitted to Google approximately ten times as often as the term "S&P 500". Moreover, the acronym "S&P" is less univocal than,

⁴For the CAC there are still 4 missing values, which we interpolate using the average of the past five observations. All missing values lie at the beginning of the sample period in August 2006, a month calm in both search queries and stock market volatility.

⁵Source: *Google Correlate* (<http://www.google.com/trends/correlate/>).

for example, “DJI”, as “S&P” is first and foremost an abbreviation for the rating agency Standard & Poor’s.

The advantage of using Google search data, in contrast to other search engines, is that Google maintains a very high market share in all countries considered. Therefore the data represent almost the entire internet searches, notably in Europe. Google’s market share is around 67.1% in the US, 91.5% in the UK, 91.2% in France and 92.7% in Germany.⁶

The data which are provided by Google are relative in nature. This means that Google does not provide the effective total number of searches, but a search volume index only. We standardize the search queries, such that the average search frequency over the sample period of 5 years equals one, allowing for an easy interpretation.

Table 4.1 also holds summary statistics for the data on search queries. Just as the realized volatility time series, the data on searches exhibit distinctive levels of skewness and kurtosis. We therefore also take logarithms of the search data (cp. Da et al. 2011). This procedure reduces both skewness and excess kurtosis, however, it is not as successful as in the case of the realized volatility. Figure 4.3 plots the autocorrelations of search queries. These are decaying fairly geometrically and much faster compared to autocorrelations of realized volatility depicted in Figure 4.2.

As already apparent from Figure 4.1, search queries and realized volatility exhibit a strong co-movement over time. The contemporary correlation of search queries and realized volatility in our sample is high and quite similar across indices. The correlation coefficients are: 0.83 (DJIA), 0.80 (FTSE), 0.80 (CAC) and 0.72 (DAX).

⁶Figures refer to June 2011. Sources: Hitwise (US), AT Internet Search Engine Barometer (Europe).

Table 4.1:
Summary statistics

This table provides descriptive statistics of realized volatility (RV) and search queries (SQ) of the DJIA, FTSE, CAC and DAX. The upper panel holds statistics for the untransformed series, the lower panel for the series after log-transformation.

	DJIA		FTSE		CAC		DAX	
	RV	SQ	RV	SQ	RV	SQ	RV	SQ
Mean	0.009	1.000	0.012	1.000	0.013	1.000	0.011	1.000
Std. Dev.	0.007	0.714	0.008	0.535	0.009	0.689	0.007	0.566
Skewness	4.01	5.54	4.07	5.97	3.79	6.11	3.01	6.81
Kurtosis	31.24	56.02	32.13	54.38	27.26	54.25	18.22	66.77
Min.	0.002	0.302	0.002	0.523	0.002	0.414	0.002	0.437
Max.	0.096	11.593	0.113	8.257	0.116	9.698	0.067	8.675
	log-RV	log-SQ	log-RV	log-SQ	log-RV	log-SQ	log-RV	log-SQ
Mean	-4.891	-0.128	-4.598	-0.067	-4.463	-0.106	-4.691	-0.069
Std. Dev.	0.568	0.453	0.525	0.318	0.517	0.406	0.508	0.318
Skewness	0.65	1.26	0.58	2.30	0.48	1.46	0.42	2.38
Kurtosis	3.71	5.67	3.87	11.26	3.96	7.96	3.51	12.88
Min.	-6.375	-1.197	-6.022	-0.648	-6.203	-0.883	-6.147	-0.829
Max.	-2.341	2.450	-2.176	2.111	-2.157	2.272	-2.699	2.160

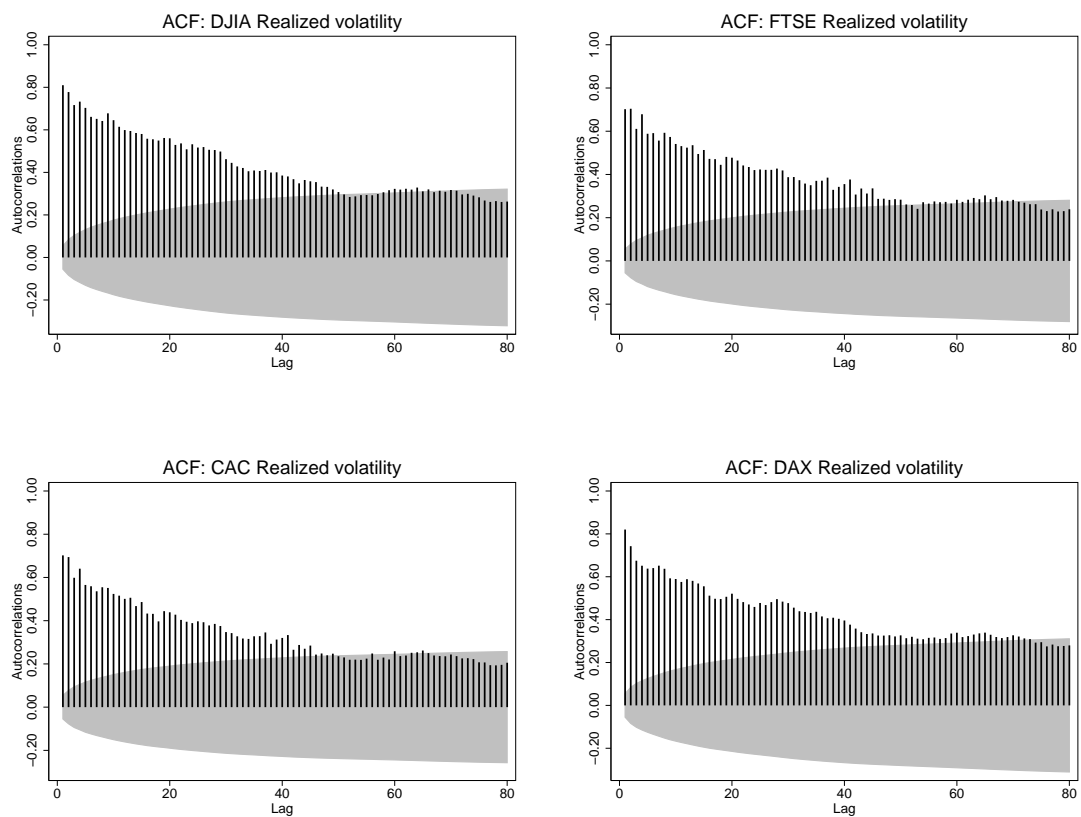


Figure 4.2:
Autocorrelations of realized volatility

This figure displays the autocorrelations of realized volatility of the stock indices DJIA, FTSE, CAC and DAX in the sample period. Shaded areas indicate 95% confidence bounds.

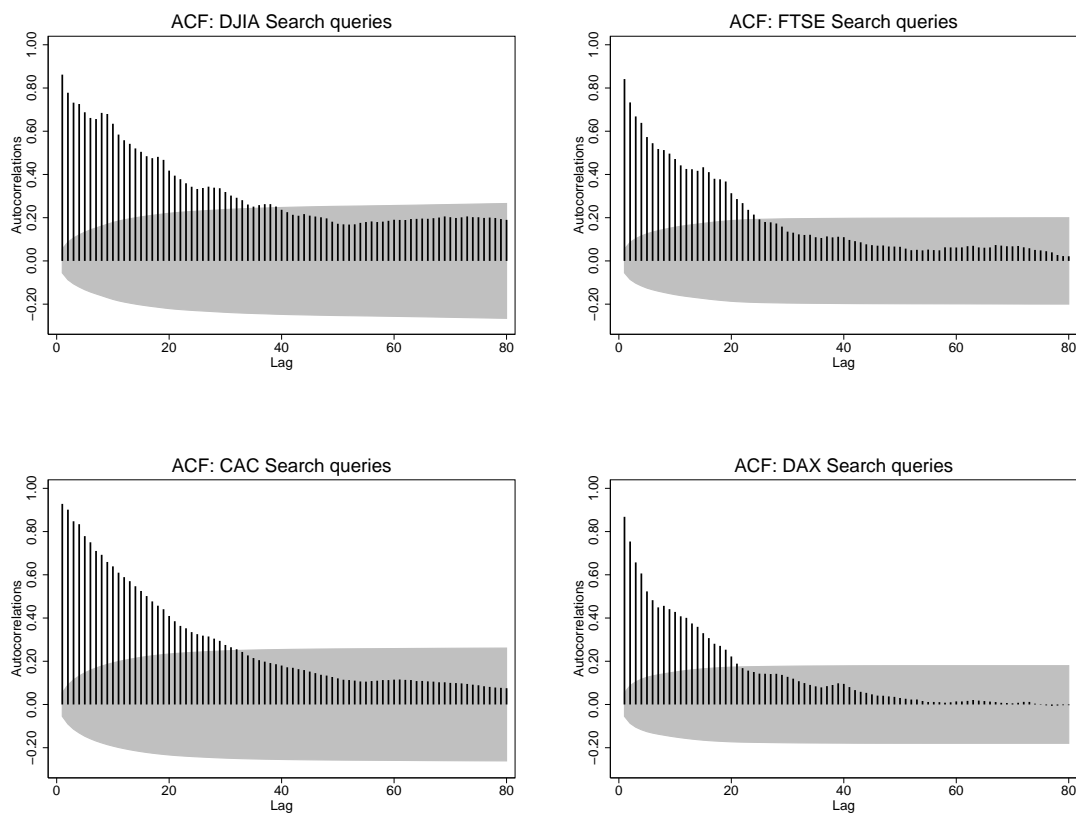


Figure 4.3:
Autocorrelations of search queries

This figure displays the autocorrelations of search queries for the stock indices DJIA, FTSE, CAC and DAX in the sample period. Shaded areas indicate 95% confidence bounds.

4.3 The dynamics of volatility and searches

4.3.1 A vector autoregressive model

In the following we study the dynamics between realized volatility and search queries. For every stock index we estimate a vector autoregressive model of order three, VAR(3), which is specified as follows:

$$\log-RV_t = c_1 + \sum_{j=1}^3 \beta_{1,j} \log-RV_{t-j} + \sum_{j=1}^3 \gamma_{1,j} \log-SQ_{t-j} + \varepsilon_{1,t} \quad (4.2a)$$

$$\log-SQ_t = c_2 + \sum_{j=1}^3 \beta_{2,j} \log-RV_{t-j} + \sum_{j=1}^3 \gamma_{2,j} \log-SQ_{t-j} + \varepsilon_{2,t}. \quad (4.2b)$$

Panel A of Table 4.2 presents the results of the four VAR models for the DJIA, FTSE, CAC and DAX. Across all indices we find significant autoregressive estimates for the realized volatility at all included lags. Search queries show significant autoregressive terms of order one, and depending on the index also significant autoregressive coefficients up to lag three.

The VAR estimation results and the Granger causality test in Panel B of Table 4.2 also reveal that in general past volatility positively influences present search queries. This effect is concentrated to the first lag $\beta_{2,1}$. One exception is the Dow Jones, where the first lag of log-SQ is slightly lower than the other indices and marginally insignificant with a p -value of 0.13. A possible explanation is that investors in the US react faster to volatility than those in Europe, which is supported by the fact that the contemporaneous correlation between searches and volatility in the US is the highest of the four countries.

The focus of our interest is how past search activity influences present volatility. For all four indices the Granger causality Lagrange multiplier (LM) test indicates that past searches provide significant information about future volatility. Past search activity influences future volatility positively and this effect is concentrated on the first lag $\gamma_{1,1}$. This coefficient is significant (on a 1% significance level) in the models of DJIA, FTSE and

DAX. In the CAC model the respective p -value is slightly above 10%, but the Granger causality LM test shows that past values of log-SQ are jointly significant.

Figure 4.4 provides the impulse response functions for one selected index, the FTSE. Impulse response functions of the other indices are alike, since the VAR estimates are very similar across indices as well. They are not reported for reasons of brevity, but available from the authors upon request.

For the calculation of impulse response functions we use a Cholesky decomposition with the economically meaningful restriction of volatility being contemporaneously exogenous, i.e. volatility can affect search queries immediately, but search queries do not contemporaneously affect volatility. The intuition behind this ordering is that there is first a fundamental volatility shock that in turn triggers retail investor attention and, thus, search queries. Search queries, on the other hand, would not rise without a preceding event on the market (see also the argumentation in Lux and Marchesi 1999).

The two top Figures present the response of log-RV and log-SQ, respectively, to a one standard-deviation shock in log-RV. As evident from the slowly decaying function, a volatility shock is highly persistent and only dies out after 30 to 40 days. The response of log-RV and log-SQ to a one standard-deviation shock in log-SQ is depicted in the two bottom figures, going from left to right. In both cases, the impact declines slightly faster than in the case of volatility shocks.

Panel C of Table 4.2 holds the long-run variance decomposition of log realized volatility and log searches. The model of Lux and Marchesi (1999) implies that volatility triggers search activity. This in turn suggests to order the variables such that volatility is contemporaneously exogenous (Ordering 1). In this case, log-RV determines a considerable amount of variance of log-SQ, ranging from 20% for the DAX to 34% for the FTSE. More importantly, the long run variance decomposition provides an answer to the question, how much of volatility can be explained by retail investors' attention. Throughout all models, the contribution of log-SQ to the variance of log-RV is significant and non-negligible: it ranges from 9% in case of the FTSE to 23% in case of the CAC. These results are in line

with Foucault et al. (2011) who document a similar order of magnitude for retail investors' contribution to volatility in stock returns.

These shares are calculated assuming that, as discussed before, volatility is contemporaneously exogenous. Of course, it is possible that retail investors react even faster to volatility shocks, i.e. at the same day, and thus contribute immediately to volatility. The above ordering does not allow for this as the respective channel is restricted. Permutating the ordering in the Cholesky decomposition, i.e. letting search queries be contemporaneously exogenous (Ordering 2 in Panel C of Table 4.2), naturally increases the contribution of log-SQ to the variance of log-RV (up to 47% in case of the DJIA). These estimates provide an upper bound for the contribution of log-SQ to the variance of log-RV. However, as outlined above the first ordering seems more appropriate and we suggest to retain the conservative lower bound as the approximate contribution of retail investors' attention to volatility.

Table 4.2:
VAR model estimation results

This table displays the estimation results of a Vector Autoregressive Model (VAR(3)) for log realized volatility (log-RV) and log search queries (log-SQ) for the indices DJIA, FTSE, CAC and DAX. Panel A provides coefficient estimates, Panel B the results of a Granger causality (LM) test and Panel C the long run forecast error variance decomposition. p -values are given in parentheses.

	Panel A: VAR estimation											
	DJIA			FTSE			CAC			DAX		
	log-RV _t	log-SQ _t	log-RV _t	log-SQ _t	log-RV _t	log-SQ _t	log-RV _t	log-SQ _t	log-RV _t	log-SQ _t	log-RV _t	log-SQ _t
log-RV _{t-1}	0.45 (0.000)	0.03 (0.132)	0.36 (0.000)	0.04 (0.015)	0.35 (0.000)	0.05 (0.000)	0.45 (0.000)	0.05 (0.000)	0.45 (0.000)	0.05 (0.000)	0.45 (0.000)	0.05 (0.000)
log-RV _{t-2}	0.21 (0.000)	0.00 (0.915)	0.26 (0.000)	0.00 (0.905)	0.25 (0.000)	0.00 (0.747)	0.17 (0.000)	0.00 (0.492)	0.17 (0.000)	-0.01 (0.492)	0.17 (0.000)	-0.01 (0.492)
log-RV _{t-3}	0.17 (0.000)	0.00 (0.868)	0.18 (0.000)	0.01 (0.502)	0.11 (0.000)	-0.03 (0.048)	0.20 (0.000)	-0.01 (0.326)	0.20 (0.000)	-0.01 (0.326)	0.20 (0.000)	-0.01 (0.326)
log-SQ _{t-1}	0.22 (0.000)	0.79 (0.000)	0.26 (0.000)	0.73 (0.000)	0.10 (0.109)	0.61 (0.000)	0.25 (0.000)	0.72 (0.000)	0.25 (0.000)	0.72 (0.000)	0.25 (0.000)	0.72 (0.000)
log-SQ _{t-2}	-0.10 (0.139)	-0.05 (0.217)	-0.17 (0.025)	0.00 (0.918)	0.03 (0.663)	0.14 (0.000)	-0.08 (0.290)	0.09 (0.013)	-0.08 (0.290)	0.09 (0.013)	-0.08 (0.290)	0.09 (0.013)
log-SQ _{t-3}	0.01 (0.925)	0.18 (0.000)	0.08 (0.180)	0.12 (0.000)	0.08 (0.237)	0.19 (0.000)	-0.04 (0.459)	0.07 (0.014)	-0.04 (0.459)	0.07 (0.014)	-0.04 (0.459)	0.07 (0.014)
Constant	-0.84 (0.000)	0.09 (0.153)	-0.93 (0.000)	0.21 (0.001)	-1.23 (0.000)	0.12 (0.037)	-0.83 (0.000)	0.13 (0.014)	-0.83 (0.000)	0.13 (0.014)	-0.83 (0.000)	0.13 (0.014)

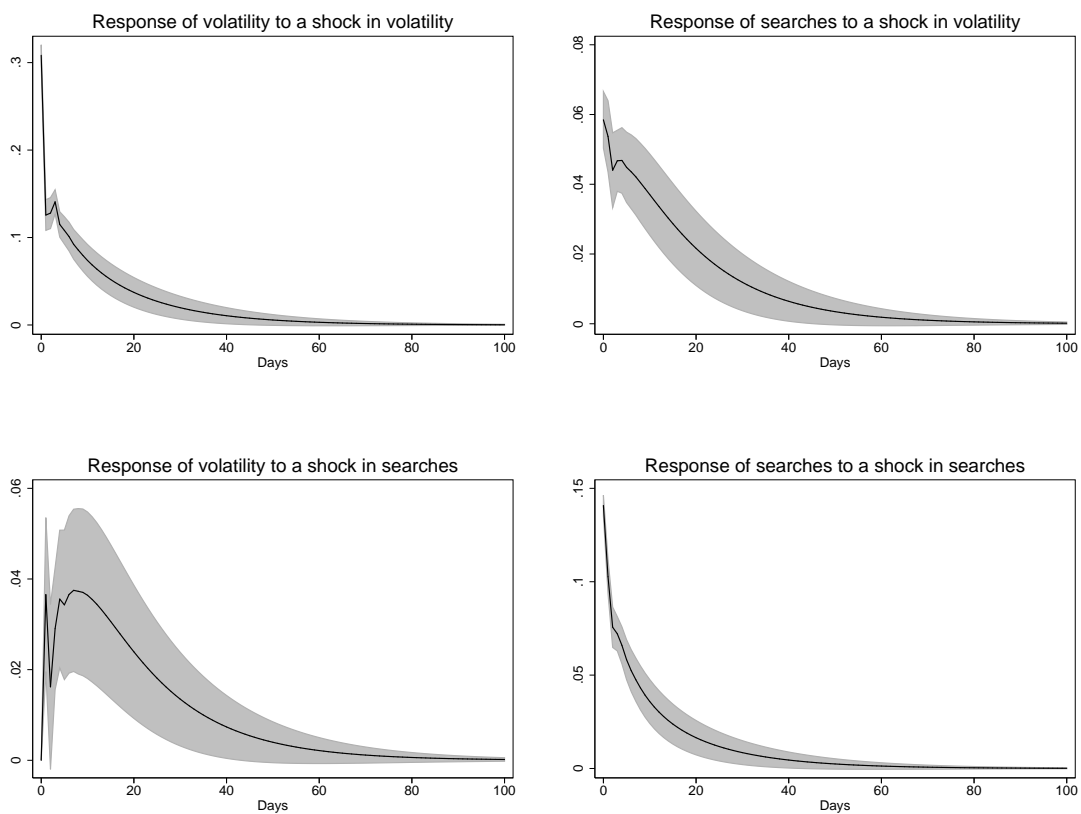


Figure 4.4:
Impulse response functions (FTSE)

The table displays the impulse response functions of the VAR(3) estimated in Table 4.2 for the FTSE. Shaded areas indicate 95% confidence bounds.

4.3.2 Do search queries add information for modeling volatility?

The key result of the VAR estimation is that search queries help to predict future volatility in addition to its own lags. One might wonder, however, whether the specific lag choice is the driver of this result. In order to rule out this explanation we turn to several other models of realized volatility. In this section we focus only on the equation of interest, the volatility equation. We use different modeling approaches which are commonly used to capture the time series properties of realized volatility and include lagged search queries in each model, testing whether searches add information. As the results of the VAR model estimation in Equation (4.2) show no significance of higher order lags we only include searches at one lag.

In particular, following Andersen, Bollerslev, Christoffersen and Diebold (2006) as well as Bollen and Inder (2002) we estimate autoregressive models with different lag length and augment these with lagged search queries $\log-SQ_{t-1}$:

$$\log-RV_t = c + \sum_{j=1}^p \beta_j \log-RV_{t-j} + \gamma_1 \log-SQ_{t-1} + \varepsilon_t. \quad (4.3)$$

We consider the lag lengths one and three. In addition to these autoregressive models we estimate Corsi's (2009) heterogeneous autoregressive (HAR) model. The HAR model has been found to capture the long-memory properties of realized volatility very well and has recently been used for example by Andersen, Bollerslev and Diebold (2007), Chen and Ghysels (2011) and Chiriac and Voev (2011). The HAR model augmented with lagged search queries reads as follows:

$$\log-RV_t = c + \beta_d \log-RV_{t-1} + \beta_w \log-RV_{t-1}^w + \beta_m \log-RV_{t-1}^m + \gamma_1 \log-SQ_{t-1} + \varepsilon_t, \quad (4.4)$$

where $\log-RV_t^w = \frac{1}{5} \sum_{j=0}^4 \log-RV_{t-j}$ and $\log-RV_t^m = \frac{1}{22} \sum_{j=0}^{21} \log-RV_{t-j}$.

As a final robustness check, we also estimate an AR(22), which includes all lags up to one month (i.e. 22 business days), in order to exclude the possibility that the aggregation of realized volatility favors the predictive power of lagged searches. This model is admittedly

Table 4.3:
Is search activity a helpful predictor of future volatility?

The table provides γ_1 coefficient estimates of lagged search queries in the univariate models described in the first column. p -values are given in parentheses.

Estimated Models:

$$\text{AR}(p): \log-RV_t = \sum_{j=1}^p \beta_j \log-RV_{t-j} + \gamma_1 \log-SQ_{t-1} + \varepsilon_t$$

$$\text{HAR}(3): \log-RV_t = \beta_d \log-RV_{t-1} + \beta_w \log-RV_{t-1}^w + \beta_m \log-RV_{t-1}^m + \gamma_1 \log-SQ_{t-1} + \varepsilon_t$$

Model:	DJIA	FTSE	CAC	DAX
AR(1)	0.22 (0.000)	0.33 (0.000)	0.36 (0.000)	0.26 (0.000)
AR(3)	0.14 (0.000)	0.18 (0.000)	0.19 (0.000)	0.16 (0.000)
HAR(3)	0.11 (0.000)	0.20 (0.000)	0.15 (0.000)	0.17 (0.000)
AR(22)	0.11 (0.000)	0.18 (0.000)	0.14 (0.000)	0.17 (0.000)

over-parameterized and not desirable from a parsimonious modeling perspective (Corsi 2009) and merely serves as a robustness check. In the forecast evaluation analysis that follows we will only consider the parsimonious model specifications.

In all four models data on the previous day's searching activity enter as an exogenous variable. We test whether γ_1 in Equations (4.3) and (4.4) is significantly different from zero to evaluate whether lagged log-SQ indeed add valuable information to the model.

Coefficient estimates of γ_1 and their corresponding p -values are presented in Table 4.3. As can be seen, lagged search queries enter significantly in all models for all indices under consideration. The findings are unambiguous and independent of the significance level as all p -values are below 1%. Even after including 22 lags of realized volatility search queries still contain significant information about future volatility. Not only the statistical but also the economic significance of lagged search queries remains. For example, the AR(3) model for the FTSE predicts that a doubling of search queries (i.e. an increase of 100%)

today increases volatility tomorrow by 18% in addition to the dynamic effects in volatility itself. The estimates of this marginal effect of lagged searches on volatility γ_1 are similar across the models AR(3), HAR(3) and AR(22). This result supports the proposition that search queries contain additional information about future volatility above and beyond the information of past volatility.

4.4 Forecast evaluation

In the following we compare the forecasting ability of the three realized volatility models AR(1), AR(3) and HAR(3) with and without search queries. We evaluate the forecasting ability of these models in- and out-of-sample as well as for multiple horizons. In order to assess the forecasting performance we consider two loss functions which are robust to possible noise in our volatility measure (see Patton 2011). These are the mean squared error (MSE) and the quasi-likelihood loss function (QL) which are defined as follows:

$$\text{MSE} = (RV_{t+1} - \widehat{RV}_{t+1|t})^2, \quad (4.5)$$

$$\text{QL} = \frac{RV_{t+1}}{\widehat{RV}_{t+1|t}} - \log \frac{RV_{t+1}}{\widehat{RV}_{t+1|t}} - 1, \quad (4.6)$$

where $\widehat{RV}_{t+1|t}$ is the respective forecast of realized volatility based upon information available up to and including time t . We also use the R^2 of a Mincer and Zarnowitz (1969) regression of the actual realized volatilities on their predicted values as follows:

$$RV_{t+1} = b_0 + b_1 \widehat{RV}_{t+1|t} + e_t. \quad (4.7)$$

Following the literature (e.g. Ait-Sahalia and Mancini 2008, Andersen et al. 2003, Ghysels, Santa-Clara and Valkanov 2006) we model log realized volatility, but evaluate the forecast by comparing realized volatility and its prediction.⁷

⁷When reversing the log transformation the forecasts are formally not optimal (Granger and Newbold 1976). However, Lütkepohl and Xu (2010) show by means of an extensive simulation study that this naïve forecast performs just as well as an optimal forecast.

4.4.1 In-sample forecast evaluation

Table 4.4 holds the results of the in-sample forecast evaluation of one-step ahead forecasts of realized volatility. The models we consider are the univariate AR(1), AR(3) and HAR(3) models and the respective augmented models including lagged search queries.

Looking only at the univariate models, we see that the AR(3) is generally better than the AR(1) and the HAR(3) is the best amongst the univariate models. These findings are in line with the literature (Corsi 2009). One exception is the CAC, where the AR(3) model seems to do reasonably well in-sample and is slightly better than the HAR(3). Comparing the univariate models (AR(1), AR(3), HAR(3)) to the SQ-augmented models (AR(1)+SQ, AR(3)+SQ, HAR(3)+SQ), we observe for all models and across all indices an improvement in performance.

Overall, the HAR model augmented with search queries, shows the best fit. Only for the CAC the AR(3) has a better (in-sample) fit than the HAR in terms of a slightly lower MSE (0.004) and a slightly higher R^2 (0.28%). However, it still holds that the model including search queries outperforms the univariate model.

4.4.2 Out-of-sample forecast evaluation

We now turn to the out-of-sample forecasts and provide 1 day, 1 week and 2 week volatility forecasts. For our initial out-of-sample forecast we estimate the models using the first two years (500 trading days) of our sample, i.e. from July 2006 to June 2008. We then re-estimate the model for every subsequent day in the sample using all past observations available, i.e. we increase the estimation window. The estimation period of the very first run ends in June 2008. Thus, we are able to compare the forecasting performance of volatility models during the almost record-high volatility of October 2008. The initial two year estimation period is still long enough and has enough variation in both volatility and search activity as to allow us to reliably estimate model parameters (compare Figure 4.1).

One-step ahead predictions can be done using the static models discussed before. For

Table 4.4:
In-sample forecast evaluation

The table compares the in-sample forecasts of the models described in the first column. AR(1), AR(3) and HAR(3) are univariate models of realized volatility only, AR(1)+SQ, AR(3)+SQ and HAR(3)+SQ are the models augmented with lagged search queries. Performance measures are the mean squared error (MSE, $\times 10^4$), the quasi-likelihood loss function (QL, $\times 10^2$) and the R^2 (in percent) of the Mincer-Zarnowitz regression. The preferred model (minimum of QL loss function and MSE, maximum of R^2) is indicated through bold numbers.

Model:	DJIA			FTSE		
	MSE	QL	R^2	MSE	QL	R^2
AR(1)	0.176	5.378	66.67	0.355	6.296	50.85
AR(1) + SQ	0.169	5.093	67.18	0.337	5.863	52.77
AR(3)	0.156	4.680	70.26	0.302	5.221	58.09
AR(3) + SQ	0.151	4.580	70.82	0.290	5.084	59.31
HAR(3)	0.149	4.503	71.47	0.293	4.990	59.23
HAR(3) + SQ	0.144	4.439	72.10	0.274	4.832	61.50

Model:	CAC			DAX		
	MSE	QL	R^2	MSE	QL	R^2
AR(1)	0.429	6.644	50.61	0.157	5.086	67.09
AR(1) + SQ	0.370	5.947	56.36	0.145	4.817	68.11
AR(3)	0.362	5.563	58.02	0.147	4.474	68.08
AR(3) + SQ	0.338	5.355	60.21	0.142	4.343	68.64
HAR(3)	0.362	5.349	57.82	0.144	4.326	68.76
HAR(3) + SQ	0.342	5.223	59.77	0.134	4.180	70.53

multi-step forecasts, however, we need to forecast log-SQ as well. For this reason we also have to model the time series properties of search queries.

Starting with the simplest model we extend the univariate AR(1) to a VAR(1) which is given as:

$$\log-RV_t = c_1 + \beta_{1,1}\log-RV_{t-1} + \gamma_{1,1}\log-SQ_{t-1} + \varepsilon_{1,t} \quad (4.8a)$$

$$\log-SQ_t = c_2 + \beta_{2,1}\log-RV_{t-1} + \gamma_{2,1}\log-SQ_{t-1} + \varepsilon_{2,t}. \quad (4.8b)$$

The model of log-SQ presented in Equation (4.8b) includes searches with one autoregressive term, but also allows for lagged log-RV to influence present log-RV. The AR(3) model is extended to a VAR(3) model the following way:

$$\log-RV_t = c_1 + \sum_{j=1}^3 \beta_{1,j}\log-RV_{t-j} + \gamma_{1,1}\log-SQ_{t-1} + \varepsilon_{1,t} \quad (4.9a)$$

$$\log-SQ_t = c_2 + \beta_{2,1}\log-RV_{t-1} + \sum_{j=1}^3 \gamma_{2,j}\log-SQ_{t-j} + \varepsilon_{2,t}. \quad (4.9b)$$

Note that the model of Equation (4.9) is a restricted version of the VAR presented earlier in Equation (4.2). Considering the results of the VAR(3) estimation in Subsection 4.3.1 we restrict the cross-influence of lagged log-RV and log-SQ on log-SQ and log-RV, respectively, to lag-order 1 in the VAR(3). That way the results are comparable to the AR(3) structure of the univariate RV-model in Subsection 4.3.2 where log-SQ entered only at lag 1 in the volatility equation (cp. Eq. (4.3)).

Finally, we augment the HAR to a Vector-HAR(3) model as follows

$$\log-RV_t = c_1 + \beta_d\log-RV_{t-1} + \beta_w\log-RV_{t-1}^w + \beta_m\log-RV_{t-1}^m + \gamma_{1,1}\log-SQ_{t-1} + \varepsilon_{1,t} \quad (4.10a)$$

$$\log-SQ_t = c_2 + \beta_{2,1}\log-RV_{t-1} + \sum_{j=1}^3 \gamma_{2,j}\log-SQ_{t-j} + \varepsilon_{2,t}. \quad (4.10b)$$

The search queries Equation (4.10b) is the same as Equation (4.9b), since we find that the

time series properties of searches are well described by three autoregressive terms and one lag of realized volatility.

We contrast the multivariate models with the univariate realized volatility models described before. That is, we compare the VAR(1) to the AR(1), the AR(3) to the VAR(3) and the HAR(3) to the VHAR(3). The univariate models AR(1), AR(3) and HAR(3) are simply equations (4.8a), (4.9a) and (4.10a) with $\gamma_{1,1}$ equal to zero. For the evaluation of weekly and biweekly forecasts of realized volatility we consider aggregated volatility over the respective time span.

Results of the out-of-sample prediction are summarized in Table 4.5. For the univariate models our results are consistent with the findings of Corsi (2009). The HAR(3) model is better at predicting realized volatility compared to the AR(3) or AR(1) model. The advantage of the HAR modeling again emerges particularly when predicting volatility at longer horizons of one or two weeks.

Turning to the multivariate models, we find that the multivariate models where searches are used as an explanatory variable always outperform the univariate, pure realized volatility models. This means that across all indices, these models have lower MSE, a lower value of the QL loss function and a higher R^2 in the Mincer-Zarnowitz regression. Adding searches is most beneficial for longer-horizon forecasts. For example in the FTSE model, the Mincer-Zarnowitz R^2 is by 3.6 percentage points higher in the multivariate VHAR(3) than in the univariate HAR(3). Also for the remaining indices, the R^2 of the VHAR(3) is by more than 3 percentage points higher compared to the HAR(3). When considering the AR-models, this difference can even be larger.

Overall, the best performing univariate model for realized volatility is the HAR model. Augmenting the HAR model with search query data further improves the forecasting performance in particular at longer horizons. What is the intuition behind this? The VHAR model benefits from modeling the dynamics of retail investors' searches and volatility and their bi-directional Granger causality. The VHAR gains from the fact that a shock in searches has a significant impact on volatility that is persistent (compare the impulse-

Table 4.5:
Out-of-sample forecast evaluation

The table compares the 1 day, 1 week and 2 weeks out-of-sample forecasts of the models described in the first column. AR(1), AR(3) and HAR(3) are univariate models of realized volatility only, VAR(1), VAR(3) and VHAR(3) are bivariate models of realized volatility (RV) and search queries (SQ). Performance measures are the mean squared error (MSE, $\times 10^4$), the quasi-likelihood loss function (QL, $\times 10^2$) and the R^2 (in percent) of the Mincer-Zarnowitz regression. The preferred model (minimum of QL loss function and MSE, maximum of R^2) is indicated through bold numbers.

Model:	1 day			1 week			2 weeks			
	MSE	QL	R^2	MSE	QL	R^2	MSE	QL	R^2	
DJIA										
AR(1)	RV	0.258	5.436	65.14	7.279	6.219	63.70	37.591	9.400	52.77
VAR(1)	RV, SQ	0.241	4.807	65.43	5.145	4.756	66.59	25.842	6.662	59.16
AR(3)	RV	0.223	4.479	69.06	4.543	3.799	72.18	22.352	5.078	66.22
VAR(3)	RV, SQ	0.214	4.227	69.25	3.943	3.328	72.66	17.653	4.256	67.94
HAR(3)	RV	0.207	4.228	70.59	3.683	3.149	74.67	15.979	3.711	70.66
VHAR(3)	RV, SQ	0.204	4.067	71.09	3.555	2.932	76.17	14.929	3.346	73.78
FTSE										
AR(1)	RV	0.478	6.785	48.15	10.40	6.263	53.01	49.905	8.608	42.91
VAR(1)	RV, SQ	0.452	6.386	51.27	8.59	5.482	63.35	41.807	7.151	58.72
AR(3)	RV	0.401	5.422	56.01	6.16	3.572	66.51	27.349	4.167	63.20
VAR(3)	RV, SQ	0.391	5.339	57.19	5.72	3.448	69.08	25.099	3.988	66.83
HAR(3)	RV	0.379	5.036	58.09	5.17	2.818	69.78	20.449	3.037	67.79
VHAR(3)	RV, SQ	0.360	4.929	60.24	4.71	2.713	72.79	18.552	2.866	71.36

Table 4.5 -Continued

Model:	1 day			1 week			2 weeks			
	MSE	QL	R ²	MSE	QL	R ²	MSE	QL	R ²	
CAC										
AR(1)	RV	0.579	6.930	46.19	13.902	7.056	44.34	64.848	9.700	29.67
VAR(1)	RV, SQ	0.486	5.502	53.30	6.623	3.875	65.38	31.010	4.748	60.03
AR(3)	RV	0.472	5.423	55.32	8.219	3.849	61.82	37.735	4.815	56.54
VAR(3)	RV, SQ	0.430	4.926	57.85	6.083	2.915	67.85	25.308	3.360	63.64
HAR(3)	RV	0.449	5.013	56.42	6.524	2.962	66.23	26.096	3.355	63.51
VHAR(3)	RV, SQ	0.425	4.709	58.61	5.947	2.512	69.86	25.222	2.743	66.76
DAX										
AR(1)	RV	0.213	5.030	63.97	7.000	5.922	51.42	34.793	8.372	36.52
VAR(1)	RV, SQ	0.191	4.788	67.36	5.689	5.192	61.58	27.743	6.725	55.23
AR(3)	RV	0.183	4.164	67.23	4.271	3.511	65.25	20.345	4.434	59.54
VAR(3)	RV, SQ	0.176	4.084	68.25	3.967	3.403	67.42	18.000	4.165	64.66
HAR(3)	RV	0.168	3.899	68.90	3.236	2.724	70.72	13.231	3.024	68.11
VHAR(3)	RV, SQ	0.160	3.820	70.40	3.101	2.656	72.43	12.140	2.842	71.42

response function of Figure 4.4). Thus, searches can improve long-run predictions. Furthermore, search queries are well described by the autoregressive time-series model allowing for good predictions of searches when the system is iterated forward.

4.4.3 Out-of-sample forecast performance over time

A further and equally important aspect in the forecasting context is the question how different volatility models behave over time. In particular, it is of interest how the models perform during high volatility phases compared to calmer periods. In this context we investigate in which phases internet search queries improve volatility forecasts. In order to do this we compare the best univariate model, the HAR(3) model, to the best bi-variate model including search activity, the VHAR(3) model.

To evaluate the gains of including search queries into the volatility model, we calculate the cumulative net sum of squared prediction errors (Net-SSE) over time. The Net-SSE compares the difference between squared prediction errors of two models. This concept was introduced by Goyal and Welch (2003) and recently used to evaluate volatility forecasts by Christiansen, Schmeling and Schrimpf (2011). The Net-SSE at time τ is given by:

$$\text{Net-SSE}(\tau) = \sum_{t=1}^{\tau} (\hat{e}_{HAR,t}^2 - \hat{e}_{VHAR,t}^2), \quad (4.11)$$

where $\hat{e}_{HAR,t}^2$ is the squared prediction error of the benchmark HAR(3) model, and $\hat{e}_{VHAR,t}^2$ is the squared prediction error of the model of interest, the VHAR(3). If the Net-SSE is positive, the VHAR(3) outperforms the benchmark HAR(3) model.

Figure 4.5 displays the Net-SSE over the out-of-sample period (July 2008 - June 2011) for all indices. The first thing to note is that for all indices and over the whole out-of-sample period the Net-SSE is positive, i.e. the VHAR with search queries outperforms the univariate HAR. This, of course, is equivalent to the results of Table 4.5, where the 1-day ahead prediction MSE of the VHAR model is smaller than that of the HAR model




Figure 4.5:
Out-of-sample performance over time

The graph shows the time variation of the out-of sample forecast measured by the cumulative sum of squared prediction error difference: $\text{Net-SSE}(\tau) = \sum_{t=1}^{\tau} (\hat{e}_{HAR,t}^2 - \hat{e}_{VHAR,t}^2)$. If the Net-SSE is positive, the model including internet searches outperforms the benchmark HAR(3) model. An increasing slope of the graph represents a better forecast performance of the VHAR(3) model (including internet searches) at this particular point in time.

throughout all indices. Thus, the overall cumulative Net-SSE corresponds to the difference in MSE between the VHAR and HAR model presented in Table 4.5.

We now turn to the question in which periods search queries add an improvement in volatility forecasts. A better forecast performance at a particular point in time is represented by an increase in the slope of the Net-SSE graph. For all four indices there is a sharp surge in Net-SSE during the high volatility phase starting in October 2008. For the DJIA there is a slight reversal during that phase, but overall there are prediction gains in this high volatility phase. When comparing Figure 4.5 to the realized volatilities of Figure 4.1 additional (smaller) rises in Net-SSE can be associated with increases in volatility. Thus, the gains of the search query data model mainly originate from turbulent times.

Figure 4.6 gives a detailed look at the volatility forecast during the financial crisis of 2008. It shows daily realized volatilities (dashed lines) for the four indices along with one-step-ahead predictions based on the HAR(3) (solid gray line) and the VHAR(3) models (solid black line) over the second half of 2008.

The plots start in July 2008, slightly before the huge increase in volatility. As can be seen, until September 2008, predictions based on the HAR(3) and the VHAR(3) models are very similar. During this calm period both models perform equally well. The advantage of using search queries in predicting realized volatility becomes apparent when volatility surges, i.e. after August 2008. We find that the univariate HAR(3) model often underestimates volatility. Furthermore, the model seems to take longer until it can finally capture the change in the realized volatility dynamics. If the model includes search queries, the predictions are closer to the actual volatility. This is particularly the case for the turbulent period of October 2008 where the VHAR(3) is clearly better able to predict the spikes in volatility than the pure HAR(3) model.

The cascading structure of the HAR(3) model seems to capture the long-memory properties or realized volatility very well. However, in a crisis period retail investors' attention is an important component and predictor of volatility. If we interpret the HAR model as

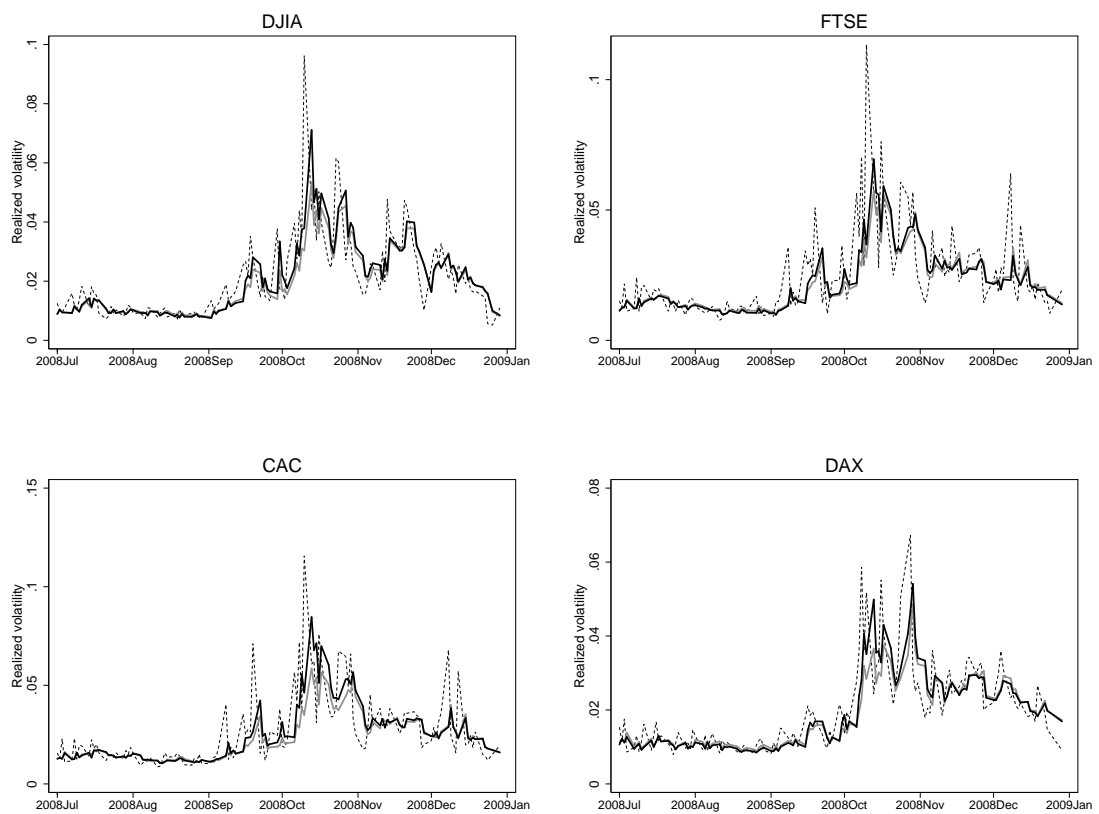


Figure 4.6:
Stock market volatility during the financial crisis

These graphs depict the realized volatilities along with predictions in the second half of 2008. The dashed lines are the realized volatilities, the solid gray lines are the out-of-sample one step ahead predictions of an HAR(3) model, the solid black line the prediction of a VHAR(3) model including search queries.

a model of agents with different time horizons (namely daily, weekly and monthly), we can understand retail investors as a fourth investor group that adds to volatility in very turbulent times.

4.5 Concluding remarks

Internet search data can describe the interest of individuals (Choi and Varian 2009a, Da et al. 2011). In this paper we use daily search query data to measure the individuals' interest in the aggregate stock market. We find that investors' attention to the stock market rises in times of high market movements. Moreover, a rise in investors' attention is followed by higher volatility. These findings are consistent with agent-based models of volatility (Lux and Marchesi 1999, Alfarano and Lux 2007).

Exploiting the fact that search queries Granger-cause volatility, we incorporate searches in several prediction models for realized volatility. Augmenting these models with search queries leads to more precise in- and out-of-sample forecasts, in particular in the long run and in high volatility phases.

Thus, search queries constitute a valuable source of information for future volatility which could essentially be used in real time. Up to now, *Google Trends* publishes search volume with a lag of only one day. Thus, long-run volatility predictions can already be improved using search query data. In principle, it would be possible to publish search volume even faster, as Google publishes the search volume for the fastest rising searches in the US through *Google Hot Trends* with only a few hours delay.⁸

⁸Google Hot Trends: <http://www.google.com/trends/hottrends>

Chapter 5

Summary and Conclusion

This thesis consists of three essays in empirical finance covering various aspects of asset prices and their relation to the behavior of investors.

Chapter 2 looks at the cross-section of stock returns and the size and value premium in the context of technology risk. We find that the risk of creative destruction plays an important role in the stock market and is priced. The growth of patent issues, patent activity growth (PAG), serves as a measure for technology shocks and creative destruction risk. While small value firms have a negative exposure to patent activity growth, large growth firms have a positive exposure to this factor. This results in an economically meaningful risk premium which can account for the size and the value premium.

The effects of technological change on asset prices have received growing attention in recent years (e.g. Nicholas 2008, Comin et al. 2009, Hsu 2009, Pástor and Veronesi 2009). A detailed analysis of how technological uncertainty affects the financial market provides an interesting agenda for future research.

Chapter 3 investigates the time-varying equity premium in the context of investor behavior. The focus of this chapter lies on one specific investor group, mutual fund investors. The key result is that mutual fund investors buy stocks when predictive variables signal high expected returns and sell stocks when predictive variables signal low expected returns. This portfolio adjustment suggests that mutual fund investors are less willing to

hold equity in bad times than the average investor. Possible explanations for this behavior are that mutual fund investors have a higher risk aversion or a higher exposure to income shocks.

Understanding the risk attitudes and portfolio choices of different investor groups is of importance for portfolio theory as pointed out by John H. Cochrane in his presidential address to the American Finance Association (AFA) in 2011: “We cannot all time the market [...]. No portfolio advice other than ‘hold the market’ can apply to everyone. A useful and durable portfolio theory must be consistent with this theorem. Our discount-rate facts and theories suggest one, built on *differences* between people” (Cochrane 2011, p. 1081, emphasis in original). By investigating the portfolio adjustment of mutual fund investors, Chapter 3 is a first step in understanding the differences between market participants. If mutual fund investors sell stocks in bad times, a question that naturally arises is, who buys these stocks and bears this risk in bad times. Answering this question is a promising avenue for further research and will help us to better understand how risk is shared in the financial market.

Chapter 4 analyzes the behavior of retail investors in the context of stock market volatility. The interest of retail investors in the stock market is measured by their search queries for index names such as Dow, FTSE, CAC, or DAX. Searches of index names and the volatility of the index show a strong co-movement over time. Furthermore, we find that searches help to predict volatility in addition to the time-series dynamics of volatility. These results can be interpreted in the context of agent-based models of volatility, where noise traders induce additional volatility. We utilize the fact that searches Granger cause volatility and include search query data into models of realized volatility. The inclusion of searches improves in-sample as well as out-of-sample forecasts.

The use of internet search queries as a measure of peoples’ attention and interest is only in its infancy, but has already found various interesting applications in economics and finance (e.g. Choi and Varian 2009a,b, Da et al. 2010a,b, 2011, Drake et al. 2011). While autoregressive time series models do very well in extrapolating trends, Choi and Varian

(2009b) argue that search query data might help to predict turning points in the data. Chapter 4 supports this by showing that times series models of realized volatility do very well in calm times, but search queries improve forecasts considerably in turbulent times.

Bibliography

- Aghion, P. and Howitt, P.: 1992, A Model of Growth Through Creative Destruction, *Econometrica* **60**(2), 323–351.
- Aït-Sahalia, Y. and Mancini, L.: 2008, Out of sample forecasts of quadratic variation, *Journal of Econometrics* **147**(1), 17–33.
- Alfarano, S. and Lux, T.: 2007, A noise trader model as a generator of apparent financial power laws and long memory, *Macroeconomic Dynamics* **11**(Supplement S1), 80–101.
- Andersen, T. G. and Bollerslev, T.: 1997, Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns, *The Journal of Finance* **52**(3), 975–1005.
- Andersen, T. G., Bollerslev, T., Christoffersen, P. F. and Diebold, F. X.: 2006, Practical Volatility and Correlation Modeling for Financial Market Risk Management, *in* M. Carey and R. M. Stulz (eds), *The Risks of Financial Institutions*, University of Chicago Press, Chicago, Illinois, chapter 17, pp. 513–548.
- Andersen, T. G., Bollerslev, T. and Diebold, F. X.: 2007, Roughing It Up: Including Jump Components in the Measurement, Modeling, and Forecasting of Return Volatility, *The Review of Economics and Statistics* **89**(4), 701–720.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. and Ebens, H.: 2001, The distribution of realized stock return volatility, *Journal of Financial Economics* **61**(1), 43–76.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P.: 2003, Modeling and Forecasting Realized Volatility, *Econometrica* **71**(2), 529–626.
- Andersen, T. G., Bollerslev, T. and Meddahi, N.: 2011, Realized Volatility Forecasting and Market Microstructure Noise, *Journal of Econometrics* **160**, 220–234.
- Bank, M., Larch, M. and Peter, G.: 2011, Google search volume and its influence on liquidity and returns of German stocks, *Financial Markets and Portfolio Management* **25**, 239–264.

- Barberis, N.: 2000, Investing for the Long Run When Returns Are Predictable, *The Journal of Finance* **55**(1), 225–264.
- Barberis, N. and Thaler, R.: 2003, A survey of behavioral finance, in M. H. G.M. Constantinides and R. Stulz (eds), *Financial Markets and Asset Pricing*, Vol. 1, Part B of *Handbook of the Economics of Finance*, Elsevier, chapter 18, pp. 1053 – 1128.
- Barro, R. J.: 1990, The Stock Market and Investment, *The Review of Financial Studies* **3**(1), 115–131.
- Ben-Rephael, A., Kandel, S. and Wohl, A.: 2010, Measuring Investor Sentiment with Mutual Fund Flows, *Journal of Financial Economics* . forthcoming.
- Bollen, B. and Inder, B.: 2002, Estimating daily volatility in financial markets utilizing intraday data, *Journal of Empirical Finance* **9**, 551–562.
- Breeden, D. T., Gibbons, M. R. and Litzenberger, R. H.: 1989, Empirical Test of the Consumption-Oriented CAPM, *The Journal of Finance* **44**(2), 231–262.
- Brennan, M. J., Schwartz, E. S. and Lagnado, R.: 1997, Strategic asset allocation, *Journal of Economic Dynamics and Control* **21**(8-9), 1377–1403. Computational financial modelling.
- Brennan, M. J., Wang, A. W. and Xia, Y.: 2004, Estimation and Test of a Simple Model of Intertemporal Capital Asset Pricing, *The Journal of Finance* **59**(4), 1743–1775.
- Bresnahan, T. F. and Trajtenberg, M.: 1995, General purpose technologies ‘Engines of growth’?, *Journal of Econometrics* **65**(1), 83–108.
- Campbell, J. Y.: 1987, Stock returns and the term structure, *Journal of Financial Economics* **18**(2), 373–399.
- Campbell, J. Y.: 1991, A Variance Decomposition for Stock Returns, *The Economic Journal* **101**(405), 157–179.
- Campbell, J. Y.: 1993, Intertemporal Asset Pricing without Consumption Data, *The American Economic Review* **83**(3), 487–512.
- Campbell, J. Y.: 1996, Understanding Risk and Return, *The Journal of Political Economy* **104**(2), 298–345.
- Campbell, J. Y. and Shiller, R. J.: 1988, The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors, *The Review of Financial Studies* **1**(3), 195–228.

- Campbell, J. Y. and Viceira, L. M.: 1999, Consumption and Portfolio Decisions When Expected Returns are Time Varying, *Quarterly Journal of Economics* **114**(2), 433–495.
- Campbell, J. Y. and Viceira, L. M.: 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press.
- Campbell, J. Y. and Vuolteenaho, T.: 2004, Bad Beta, Good Beta, *The American Economic Review* **94**(5), 1249–1275.
- Carhart, M.: 1997, On Persistence in Mutual Fund Performance, *The Journal of Finance* **52**, 57–82.
- Chalmers, J., Kaul, A. and Phillips, B.: 2010, Economic conditions, flight to quality and mutual fund flows, *Unpublished Working Paper*.
- Chan, K. C. and Chen, N.-F.: 1991, Structural and Return Characteristics of Small and Large Firms, *The Journal of Finance* **46**(4), 1467–1484.
- Chan, K. C., Chen, N.-f. and Hsieh, D. A.: 1985, An exploratory investigation of the firm size effect, *Journal of Financial Economics* **14**(3), 451–471.
- Chan, Y. L. and Kogan, L.: 2002, Catching up with the Joneses: Heterogeneous Preferences and the Dynamics of Asset Prices, *The Journal of Political Economy* **110**(6), 1255–1285.
- Chen, L., Petkova, R. and Zhang, L.: 2008, The expected value premium, *Journal of Financial Economics* **87**(2), 269 – 280.
- Chen, N.-F.: 1991, Financial Investment Opportunities and the Macroeconomy, *The Journal of Finance* **46**(2), 529–554.
- Chen, N.-F., Roll, R. and Ross, S. A.: 1986, Economic Forces and the Stock Market, *The Journal of Business* **59**(3), 383–403.
- Chen, X. and Ghysels, E.: 2011, News-Good or Bad-and Its Impact on Volatility Predictions over Multiple Horizons, *Review of Financial Studies* **24**(1), 46–81.
- Chiriac, R. and Voev, V.: 2011, Modelling and forecasting multivariate realized volatility, *Journal of Applied Econometrics* **26**(6), 922–947.
- Choi, H. and Varian, H.: 2009a, Predicting initial claims for unemployment benefits, *Working Paper*.
- Choi, H. and Varian, H.: 2009b, Predicting the present with Google trends, *Working Paper* pp. 1–23.

- Christiansen, C., Schmeling, M. and Schrimpf, A.: 2011, A Comprehensive Look at Financial Volatility Prediction by Economic Variables, *CREATES Research Papers* .
- Cochrane, J.: 1999, New Facts in Finance, *Economic Perspectives* **23**(3), 36–38.
- Cochrane, J.: 2005, *Asset Pricing*, 2 edn, Princeton University Press, Princeton, NJ.
- Cochrane, J.: 2008, Financial Markets and the Real Economy, *Handbook of the Equity Risk Premium*, Elsevier.
- Cochrane, J. H.: 1994, Permanent and Transitory Components of GNP and Stock Prices, *The Quarterly Journal of Economics* **109**(1), 241–265.
- Cochrane, J. H.: 2007, Portfolio theory, *Working Paper, University of Chicago* .
- Cochrane, J. H.: 2011, Presidential Address: Discount Rates, *The Journal of Finance* **66**(4), 1047–1108.
- Comin, D., Gertler, M. and Santacreu, A. M.: 2009, Technology Innovation and Diffusion as Sources of Output and Asset Price Fluctuations, *Working Paper* .
- Constantinides, G. M. and Duffie, D.: 1996, Asset Pricing with Heterogeneous Consumers, *The Journal of Political Economy* **104**(2), 219–240.
- Corsi, F.: 2009, A Simple Approximate Long-Memory Model of Realized Volatility, *Journal of Financial Econometrics* **7**(2), 174–196.
- Da, Z., Engelberg, J. and Gao, P.: 2010a, In search of earnings predictability, *Working Paper* .
- Da, Z., Engelberg, J. and Gao, P.: 2010b, The Sum of All FEARS: Investor Sentiment and Asset Prices, *Working Paper* .
- Da, Z., Engelberg, J. and Gao, P.: 2011, In Search of Attention, *The Journal of Finance* **66**(5), 1461–1499.
- D’Avolio, G.: 2002, The market for borrowing stock, *Journal of Financial Economics* **66**(2-3), 271–306.
- Dimpfl, T. and Jank, S.: 2011, Can internet search queries help to predict stock market volatility?, *Working Paper* .
- Drake, M., Roulstone, D. and Thornock, J.: 2011, Investor Information Demand: Evidence from Google Searches around Earnings Announcements, *Working Paper* .

- Dumas, B.: 1989, Two-Person Dynamic Equilibrium in the Capital Market, *The Review of Financial Studies* **2**(2), 157–188.
- Edelen, R. M. and Warner, J. B.: 2001, Aggregate price effects of institutional trading: a study of mutual fund flow and market returns, *Journal of Financial Economics* **59**(2), 195–220.
- Edwards, F. R. and Zhang, X.: 1998, Mutual Funds and Stock and Bond Market Stability, *Journal of Financial Services Research* **13**, 257–282.
- Fama, E. F.: 1970a, Efficient Capital Markets: A Review of Theory and Empirical Work, *The Journal of Finance* **25**(2), pp. 383–417.
- Fama, E. F.: 1970b, Multiperiod Consumption-Investment Decisions, *The American Economic Review* **60**(1), 163–174.
- Fama, E. F.: 1981, Stock Returns, Real Activity, Inflation, and Money, *The American Economic Review* **71**(4), 545–565.
- Fama, E. F.: 1990, Stock Returns, Expected Returns, and Real Activity, *The Journal of Finance* **45**(4), 1089–1108.
- Fama, E. F.: 1991, Efficient Capital Markets: II, *The Journal of Finance* **46**(5), 1575–1617.
- Fama, E. F. and French, K. R.: 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* **22**(1), 3–25.
- Fama, E. F. and French, K. R.: 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* **25**(1), 23–49.
- Fama, E. F. and French, K. R.: 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* **33**(1), 3–56.
- Fama, E. F. and French, K. R.: 1995, Size and Book-to-Market Factors in Earnings and Returns, *The Journal of Finance* **50**(1), 131–155.
- Fama, E. F. and French, K. R.: 1996, Multifactor Explanations of Asset Pricing Anomalies, *The Journal of Finance* **51**(1), 55–84.
- Fama, E. F. and French, K. R.: 2002, The Equity Premium, *The Journal of Finance* **57**(2), 637–659.
- Fant, L. F.: 1999, Investment behavior of mutual fund shareholders: The evidence from aggregate fund flows, *Journal of Financial Markets* **2**(4), 391–402.

- Ferson, W. E. and Harvey, C. R.: 1991, The Variation of Economic Risk Premiums, *The Journal of Political Economy* **99**(2), 385–415.
- Foucault, T., Sraer, D. and Thesmar, D. J.: 2011, Individual Investors and Volatility, *The Journal of Finance* **66**(4), 1369–1406.
- Friesen, G. C. and Sapp, T. R.: 2007, Mutual fund flows and investor returns: An empirical examination of fund investor timing ability, *Journal of Banking & Finance* **31**(9), 2796–2816.
- Frisch, R. and Waugh, F. V.: 1933, Partial Time Regressions as Compared with Individual Trends, *Econometrica* **1**(4), 387–401.
- Geske, R. and Roll, R.: 1983, The Fiscal and Monetary Linkage Between Stock Returns and Inflation, *The Journal of Finance* **38**(1), 1–33.
- Ghysels, E., Santa-Clara, P. and Valkanov, R.: 2006, Predicting Volatility: Getting the Most out of Return Data Sampled at Different Frequencies, *Journal of Econometrics* **1-2**, 59–95.
- Ghysels, E. and Sinko, A.: 2011, Volatility Forecasting and Microstructure Noise, *Journal of Econometrics* **160**, 257–271.
- Ginsberg, J., Mohebbi, M. H., Patel, R. S., Brammer, L., Smolinski, M. S. and Brilliant, L.: 2009, Detecting influenza epidemics using search engine query data, *Nature* **457**(7232), 1012–1014.
- Goyal, A. and Welch, I.: 2003, Predicting the Equity Premium with Dividend Ratios, *Management Science* **49**(5), 639–654.
- Grammig, J. and Jank, S.: 2010, Creative Destruction and Asset Prices, *Working Paper*.
- Granger, C. W. J.: 1969, Investigating Causal Relations by Econometric Models and Cross-spectral Methods, *Econometrica* **37**(3), 424–438.
- Granger, C. W. J. and Newbold, P.: 1976, Forecasting Transformed Series, *Journal of the Royal Statistical Society. Series B (Methodological)* **38**(2), 189–203.
- Grossman, G. M. and Helpman, E.: 1991, Quality Ladders in the Theory of Growth, *The Review of Economic Studies* **58**(1), 43–61.
- Grossman, S. J. and Zhou, Z.: 1996, Equilibrium Analysis of Portfolio Insurance, *The Journal of Finance* **51**(4), 1379–1403.

- Helpman, E. and Trajtenberg, M.: 1994, A Time to Sow and a Time to Reap: Growth Based on General Purpose Technologies, *NBER Working Paper* (4854).
- Hodrick, R. J.: 1992, Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement, *The Review of Financial Studies* **5**(3), 357–386.
- Hsu, P.-H.: 2009, Technological innovations and aggregate risk premiums, *Journal of Financial Economics* **94**(2), 264–279.
- Jacobs, H. and Weber, M.: forthcoming, The Trading Volume Impact of Local Bias: Evidence from a Natural Experiment, *Review of Finance* .
- Jagannathan, R. and Wang, Z.: 1996, The Conditional CAPM and the Cross-Section of Expected Returns, *The Journal of Finance* **51**(1), 3–53.
- Jank, S.: 2011, Mutual fund flows, expected returns, and the real economy, *Working Paper* .
- Jegadeesh, N. and Titman, S.: 1993, Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency, *The Journal of Finance* **48**(1), 65–91.
- Kaul, G.: 1987, Stock returns and inflation : The role of the monetary sector, *Journal of Financial Economics* **18**(2), 253–276.
- Keswani, A. and Stolin, D.: 2008, Which Money Is Smart? Mutual Fund Buys and Sells of Individual and Institutional Investors, *The Journal of Finance* **63**(1), 85–118.
- Lakonishok, J., Shleifer, A. and Vishny, R. W.: 1994, Contrarian Investment, Extrapolation, and Risk, *The Journal of Finance* **49**(5), 1541–1578.
- Lamont, O. A.: 2001, Economic tracking portfolios, *Journal of Econometrics* **105**(1), 161–184.
- Lamont, O. and Thaler, R.: 2003, Can the Market Add and Subtract? Mispricing in Tech Stock Carve-outs, *Journal of Political Economy* **111**(2), 227–268.
- Lemmon, M. and Portniaguina, E.: 2006, Consumer Confidence and Asset Prices: Some Empirical Evidence, *The Review of Financial Studies* **19**(4), 1499–1529.
- Lettau, M. and Ludvigson, S.: 2001, Consumption, Aggregate Wealth, and Expected Stock Returns, *The Journal of Finance* **56**(3), 815–849.

- Lettau, M. and Ludvigson, S. C.: 2004, Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption, *The American Economic Review* **94**(1), 276–299.
- Lettau, M. and Ludvigson, S. C.: 2005, Expected returns and expected dividend growth, *Journal of Financial Economics* **76**(3), 583–626.
- Lewellen, J., Nagel, S. and Shanken, J.: 2010, A skeptical appraisal of asset pricing tests, *Journal of Financial Economics* **96**(2), 175–194.
- Liew, J. and Vassalou, M.: 2000, Can book-to-market, size and momentum be risk factors that predict economic growth?, *Journal of Financial Economics* **57**(2), 221–245.
- Lintner, J.: 1965, The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets, *The Review of Economics and Statistics* **47**(1), 13–37.
- Ljung, G. M. and Box, G. E. P.: 1978, On a Measure of Lack of Fit in Time Series Models, *Biometrika* **65**(2), 297–303.
- Ludvigson, S. C.: 2004, Consumer Confidence and Consumer Spending, *The Journal of Economic Perspectives* **18**(2), 29–50.
- Lütkepohl, H. and Xu, F.: 2010, The role of the log transformation in forecasting economic variables, *Empirical Economics* pp. 1–20.
- Lux, T. and Marchesi, M.: 1999, Scaling and criticality in a stochastic multi-agent model of a financial market, *Nature* **397**(6719), 498–500.
- Mankiw, N. G.: 1986, The equity premium and the concentration of aggregate shocks, *Journal of Financial Economics* **17**(1), 211–219.
- Merton, R. C.: 1973, An Intertemporal Capital Asset Pricing Model, *Econometrica* **41**(5), 867–887.
- Mincer, J. A. and Zarnowitz, V.: 1969, The Evaluation of Economic Forecasts, in J. A. Mincer (ed.), *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*, Studies in Business Cycles, NBER.
- Mitchell, M., Pulvino, T. and Stafford, E.: 2002, Limited Arbitrage in Equity Markets, *The Journal of Finance* **57**(2), 551–584.
- Mossin, J.: 1966, Equilibrium in a Capital Asset Market, *Econometrica* **34**(4), 768–783.

- Nesbitt, S.: 1995, Buy High, Sell Low: Timing Errors in Mutual Fund Allocations, *The Journal of Portfolio Management* **22**(1), 57–60.
- Nicholas, T.: 2008, Does innovation cause stock market runups? Evidence from the great crash, *The American Economic Review* **98**(4), 1370–1396.
- Parker, J. A. and Julliard, C.: 2005, Consumption Risk and the Cross Section of Expected Returns, *Journal of Political Economy* **113**(1), 185–222.
- Pástor, L. and Veronesi, P.: 2009, Technological Revolutions and Stock Prices, *The American Economic Review* **99**(4), 1451–1483.
- Patton, A. J.: 2011, Volatility forecast comparison using imperfect volatility proxies, *Journal of Econometrics* **160**(1), 246–256.
- Petkova, R.: 2006, Do the Fama-French Factors Proxy for Innovations in Predictive Variables?, *The Journal of Finance* **61**(2), 581–612.
- Petkova, R. and Zhang, L.: 2005, Is value riskier than growth?, *Journal of Financial Economics* **78**(1), 187 – 202.
- Rakowski, D. and Wang, X.: 2009, The dynamics of short-term mutual fund flows and returns: A time-series and cross-sectional investigation, *Journal of Banking & Finance* **33**(11), 2102–2109.
- Roll, R.: 1984, Orange Juice and Weather, *The American Economic Review* **74**(5), 861–880.
- Sapp, T. and Tiwari, A.: 2004, Does Stock Return Momentum Explain the "Smart Money" Effect?, *The Journal of Finance* **59**(6), 2605–2622.
- Schumpeter, J. A.: 1961, *Capitalism, Socialism and Democracy*, 4 edn, London : Allen & Unwin.
- Schwert, G. W.: 1990, Stock Returns and Real Activity: A Century of Evidence, *The Journal of Finance* **45**(4), 1237–1257.
- Segerstrom, P. S., Anant, T. C. A. and Dinopoulos, E.: 1990, A Schumpeterian Model of the Product Life Cycle, *The American Economic Review* **80**(5), 1077–1091.
- Shanken, J.: 1992, On the Estimation of Beta-Pricing Models, *The Review of Financial Studies* **5**(1), 1–33.
- Sharpe, W. F.: 1964, Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk, *The Journal of Finance* **19**(3), 425–442.

- Shiller, R. J., Fischer, S. and Friedman, B. M.: 1984, Stock Prices and Social Dynamics, *Brookings Papers on Economic Activity* **1984**(2), 457–510.
- Shleifer, A. and Summers, L. H.: 1990, The Noise Trader Approach to Finance, *The Journal of Economic Perspectives* **4**(2), pp. 19–33.
- Townes, C.: 2003, The first laser, in L. Garwin and T. Lincoln (eds), *A century of nature: twenty-one discoveries that changed science and the world*, The University of Chicago Press, pp. 107–112.
- Vassalou, M.: 2003, News related to future GDP growth as a risk factor in equity returns, *Journal of Financial Economics* **68**(1), 47–73.
- Vassalou, M. and Xing, Y.: 2004, Default Risk in Equity Returns, *The Journal of Finance* **59**(2), 831–868.
- Wang, J.: 1996, The term structure of interest rates in a pure exchange economy with heterogeneous investors, *Journal of Financial Economics* **41**(1), 75–110.
- Warther, V. A.: 1995, Aggregate mutual fund flows and security returns, *Journal of Financial Economics* **39**(2-3), 209–235.
- Yogo, M.: 2006, A Consumption-Based Explanation of Expected Stock Returns, *The Journal of Finance* **61**(2), 539–580.
- Zhang, L.: 2005, The Value Premium, *The Journal of Finance* **60**(1), 67–103.
- Zheng, L.: 1999, Is Money Smart? A Study of Mutual Fund Investors' Fund Selection Ability, *The Journal of Finance* **54**(3), 901–933.