Risk and Return Dynamics

Citation for published version (APA):

Cosemans, M. M. J. E. (2010). *Risk and Return Dynamics*. [Doctoral Thesis, Maastricht University]. Datawyse / Universitaire Pers Maastricht. <https://doi.org/10.26481/dis.20100401mc>

Document status and date:

Published: 01/01/2010

DOI:

[10.26481/dis.20100401mc](https://doi.org/10.26481/dis.20100401mc)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Risk and Return Dynamics

Mathijs Cosemans

© Mathijs Cosemans, 2010

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form, or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission in writing from the author.

Front cover image: © SuperStock

Back cover image: © Fotolia

This book was typeset by the author using L^AT_EX.

Published by Universitaire Pers Maastricht

ISBN: 978-90-5278-923-1

Printed in the Netherlands by Datawyse Maastricht

Risk and Return Dynamics

PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit Maastricht,
op gezag van de Rector Magnificus,
Prof. mr. G.P.M.F. Mols,
volgens het besluit van het College van Decanen,
in het openbaar te verdedigen
op donderdag 1 april 2010 om 14.00 uur

door

Mathijs Maria Jacobus Elisabeth Cosemans

Promotores:

Prof. dr. R.M.M.J. Bauer
Prof. dr. P.M.A. Eichholtz
Prof. dr. P.C. Schotman

Beoordelingscommissie:

Prof. dr. ir. J.M.E. Pennings (voorzitter)
Prof. dr. J.J.A.G. Driessen (Universiteit van Tilburg)
Prof. dr. J.R.Y.J. Urbain

Dit onderzoek is financieel mogelijk gemaakt door de Maastricht Research School of Economics of Technology and Organizations (METEOR) en Inquire Europe.

Acknowledgements

After writing two Master's theses I was convinced that I liked doing empirical research in finance very much. Consequently, I decided to take the next step and started working on my Ph.D. thesis in September 2005. The subject of my dissertation, risk and return dynamics, not only applies to financial markets but also relates to academic research in a broader sense. After all, writing a thesis is a dynamic process that requires an investment of four years with a return that is uncertain at the outset and dependent on the interaction of multiple factors. It has been a great experience and I am very pleased with the outcome: this dissertation. Therefore, I would like to express my gratitude to some of the “factors” that have contributed to my research in various ways.

First, I would like to thank my supervisors, Rob Bauer, Piet Eichholtz, and Peter Schotman, for giving me the opportunity to pursue a Ph.D. degree and for their guidance and support. Rob, I am grateful for your confidence in me and for your insights from practice. Piet, I thank you for encouraging me and for giving me the freedom to pursue my own research interests. Peter, I am indebted to you for your constructive comments, for teaching me to critically reflect on my work, and for always having a spare moment to talk about research. I also thank the members of the assessment committee, Joost Driessen, Joost Pennings, and Jean-Pierre Urbain, for their careful reading of the manuscript.

Another word of thanks goes to Luis Viceira for giving me the opportunity to spend a semester at Harvard University as a visiting research fellow. I truly enjoyed my time in Boston and learned a lot from the courses and seminars at Harvard. I also would like to express my gratitude to the Finance Department at Maastricht University for the pleasant working atmosphere and for the financial support that made it possible for me to present my research at conferences across Europe and the United States. I am grateful to Frank Lutgens, Paul Smeets, and Robin Braun for all the research- and non-research-related conversations we had.

Special thanks to Rik Frehen, who first became my roommate and later a co-author and close friend. You not only improved the quality of my research but also contributed to my happiness at work. I even look back with a smile at our extreme experiences, like almost crashing into a coyote while crossing Death Valley.

ACKNOWLEDGEMENTS

Mijn grootste dank gaat uit naar mijn familie. Judith, jij hebt als “grote” zus altijd het goede voorbeeld gegeven en mij gemotiveerd door zelf te promoveren. Het enige verschil is dat ik graag cellen in een spreadsheet bestudeer terwijl jij liever cellen onder de microscoop bekijkt. Mijn ouders Albert en Gerda ben ik dankbaar voor hun onvoorwaardelijke steun, liefde en interesse gedurende mijn hele leven. Zonder jullie zorg en aandacht was ik nooit zo ver gekomen, niet alleen als onderzoeker, maar vooral als mens.

Mathijs Cosemans
Born, November 2009

Contents

1	Introduction	1
1.1	Risk and Return Dynamics	3
1.2	Investor Irrationality in Financial Markets	4
1.3	Aim and Outline of the Thesis	6
2	Conditional Asset Pricing and Stock Market Anomalies in Europe	9
2.1	Introduction	9
2.2	Conditional Asset Pricing: Theory and Evidence	12
2.3	Methodology	13
2.3.1	Time Series Test Methodology	13
2.3.2	Cross-Sectional Framework	14
2.4	Data and Descriptive Statistics	16
2.5	Empirical Results	20
2.5.1	Time Series Evidence on the Unconditional Three-Factor Model	20
2.5.2	Predictability of Size and Book-to-Market Portfolio Returns	21
2.5.3	Time-Varying Betas in the Three-Factor Model	25
2.5.4	Conditional Alphas in the Three-Factor Model	25
2.5.5	Cross-Sectional Evidence on the Three-Factor Model	27
2.5.6	Business Cycle Risk and Momentum	29
2.5.7	Robustness of Empirical Results	31
2.6	Conclusion	32
3	Efficient Estimation of Firm-Specific Betas and its Benefits for Asset Pricing Tests and Portfolio Choice	35
3.1	Introduction	35
3.2	The Model	39
3.3	Methodology	41
3.3.1	Bayesian Methods	41
3.3.2	Prior Distributions	42

CONTENTS

3.3.3	Bayesian Inference	43
3.4	Data	43
3.5	Empirical Results	44
3.5.1	Beta Specification	45
3.5.2	Beta Estimation	51
3.5.3	Portfolio Heterogeneity	53
3.6	Applications of Firm-Specific Betas	56
3.6.1	Cross-Sectional Tests of the Conditional CAPM	58
3.6.2	Beta Forecasts and Minimum Variance Portfolios	60
3.7	Conclusion	64
3.A	Appendix A: Posterior Distributions	66
3.A.1	Joint Posterior Distribution	66
3.A.2	Conditional Posterior Distributions	67
3.B	Appendix B: Cross-Sectional Tests	71
4	Long and Short Run Correlation Risk in Stock Returns	73
4.1	Introduction	73
4.2	Related Literature on Variance Risk and Correlation Risk	76
4.2.1	Variance and Correlation Risk Premia in Aggregate Stock Returns	76
4.2.2	Cross-Sectional Evidence on Variance and Correlation Risk	77
4.3	Time Series Analysis of Variance and Correlation Risk	78
4.3.1	Measurement of Variance and Correlation Risk Premia	78
4.3.2	Data Description	80
4.3.3	Predicting Stock Market Returns with Variance and Correlation Risk Premia	82
4.4	Long and Short Run Variances and Correlations	89
4.4.1	ICAPM with Time-Varying Variances and Correlations	89
4.4.2	Specification of Volatility and Correlation Components	90
4.4.3	Estimation of Long and Short Run Volatilities and Correlations	95
4.5	Cross-Sectional Price of Variance and Correlation Risk	100
4.6	Conclusion	109
5	Option Trading and Individual Investor Performance	111
5.1	Introduction	111
5.2	Data	114
5.2.1	Individual Investors and Financial Markets in the Netherlands	114
5.2.2	Data Set	115
5.3	Methods	116
5.3.1	Measuring Investor Performance	116
5.3.2	Performance Attribution	118
5.4	Results	120
5.4.1	Option Trading and Investor Performance	120
5.4.2	Investor Sentiment and Market Timing	125
5.4.3	Hedging and Gambling Motivations for Option Trading	131

5.4.4	Option Trading Skills	137
5.5	Conclusion	139
5.A	Appendix: Cross-Sectional Risk and Style Adjustment	142
6	Conclusion	143
6.1	Summary of Main Findings	143
6.2	Implications of Empirical Results	147
6.3	Suggestions for Future Research	148
	Bibliography	151
	Nederlandse Samenvatting	163
	Curriculum Vitae	171

Chapter 1

Introduction

The concept of risk plays a pervasive role in our lives and takes on different forms. Risk affects decisions in a wide range of areas, including business, engineering, medicine, and science.¹ Although the exact meaning of risk depends on the context, the term generally refers to the possibility of a loss or other undesirable outcome. The economist Knight (1921) makes an important distinction between risk and uncertainty. According to this interpretation, *risk* pertains to situations where the decision maker can assign mathematical probabilities to the randomness which he is faced with whereas *uncertainty* refers to situations in which this randomness cannot be expressed in terms of mathematical probabilities. Keynes (1937) illustrates this distinction by noting that the outcome of a game of roulette is risky but not uncertain. In contrast, the prospect of a war is uncertain because there is no scientific basis on which to form any calculable probability. Since risk has such a profound impact on society, proper measurement and management of risk are of crucial importance for economic growth and technological progress. These intricate tasks are further complicated by the dynamic nature of risk.

This thesis is about the measurement and pricing of time-varying risk in stock markets. While the importance of risk in games of chance was already recognized centuries ago, the concept was not applied to stock markets until Harry Markowitz developed his portfolio selection theory in the 1950s, for which he received the Nobel prize in Economics in 1990 (joint with William Sharpe and Merton Miller). Markowitz (1952) argues that investors will only hold mean-variance efficient portfolios, which means that the portfolios should maximize expected returns for a given level of risk. The key element of portfolio theory is the concept of diversification, which says that the risk of a portfolio can be reduced by allocating wealth across different stocks as long as their prices do not move together perfectly. The risk of a portfolio, measured by its return variance, therefore not only depends on the variances of the individual stocks in the portfolio, but also on the covariances between these securities. In fact, because of diversification, the portfolio variance is more dependent on these covariance terms than on the individual risks.

¹Bernstein (1996) provides a fascinating account of the history of risk in society.

The insights from portfolio theory led Sharpe (1964) and Lintner (1965) to develop a model to price financial assets, the Capital Asset Pricing Model (CAPM). They assume that investors have homogeneous expectations about the joint distribution of asset returns and can borrow and lend at a risk-free rate of interest. Under these conditions and with no market frictions, all investors face the same set of portfolio opportunities and hold the same portfolio of risky assets, which must be the market portfolio. They can combine this investment in the market portfolio with a position in the risk-free asset to achieve their desired risk exposure. If all investors hold the same portfolio of risky assets, the relevant risk of an individual security is its contribution to the risk of the market portfolio. This contribution is known as the market beta of the security and is defined as the covariance between the asset return and the market return, standardized by the variance of the market portfolio. Beta therefore measures the sensitivity of the return on an asset to variation in the return on the market. According to the CAPM, investors are only compensated for taking systematic risk, measured by beta, because stock-specific risk can be diversified away. The model therefore states that the expected return on an asset is positively and linearly related to its beta.

Early empirical tests of the CAPM supported its prediction that the cross-section of expected returns should be completely explained by beta (Black, Jensen, and Scholes (1972) and Fama and MacBeth (1973)). However, later studies discovered return patterns that are inconsistent with the implications of the CAPM, including the size effect documented by Banz (1981), the value premium discovered by Basu (1977), the mean reversion observed by DeBondt and Thaler (1985), and the momentum effect detected by Jegadeesh and Titman (1993). Another blow to the CAPM was given by Fama and French (1992), who show in an influential study that size and book-to-market equity capture the cross-sectional variation in returns associated with equity characteristics and beta. In addition, they find that the relation between market beta and average return is flat when tests control for size. In a subsequent paper, Fama and French (1993) propose a three-factor model that extends the CAPM with two factors designed to capture the size effect and value premium. Despite its ability to clear up most of the CAPM anomalies, the three-factor model is often criticized for the lack of a sound theoretical motivation for the inclusion of the size and value factors. Moreover, it does not explain the momentum effect observed in stock returns.

Many different explanations have been offered for the empirical failure of the CAPM and its multifactor generalizations. The main debate is between defenders of rational asset pricing theory who believe in a risk-based explanation for the anomalous patterns in returns and attackers of market efficiency who argue that mispricing occurs because investors exhibit irrational behavior in the interpretation of new information. Cochrane (2005) points out that ultimately, this discussion boils down to the fundamental question “whether asset pricing theory describes the way the world *does* work or the way the world *should* work.” Proponents of neoclassical finance interpret deviations from the predictions of the model as evidence that the model specification is incomplete while believers in a behavioral story attribute these anomalies to investor irrationality.

1.1 Risk and Return Dynamics

According to the risk story, a possible reason for the inability of the CAPM and other rational models to explain the cross-section of returns is the key assumption of these models that the risk exposure of assets is constant. This assumption is not very plausible given the dynamic nature of our economies and financial markets. Because the risk of a firm's cash flow is sensitive to changes in the macroeconomic environment, firm betas and expected returns will fluctuate. Betas can also vary due to changes in firm-specific conditions, for instance when a firm changes its main line of business. More generally, Fama (1965) and Samuelson (1965) point out that in an efficient market, movements in stock prices reflect revisions in the expected value of discounted dividends. These changes in forecasts are driven by new information about the prospects of the firm and the expected return used to discount future cash flows. Campbell and Shiller (1988) formalize this idea by decomposing unexpected stock returns into revisions in expected dividend growth (cash flow news) and shocks to expected returns (discount rate news). In this framework, return variance depends on the variance of cash flow news and discount rate news. Similarly, return correlations depend on the correlation between the dividend news of assets and on the correlation between their discount rate news. As the intensity of news varies over time and across firms, volatilities and correlations also change through time. Since betas depend on variances and correlations, any variation in these determinants induces fluctuations in betas.

If true betas vary over time, static asset pricing models are misspecified and give an incomplete description of stock returns. This observation has motivated the development of conditional asset pricing models in which risk loadings can change through time. These models imply that the conditional expected return on a stock is a linear function of conditional betas that measure its time-varying sensitivity to systematic risk factors. Hansen and Richard (1987) provide theoretical support for dynamic models by showing that a conditional version of the CAPM can hold perfectly even if its unconditional counterpart fails. Nevertheless, empirical evidence on the performance of conditional models is far from conclusive.

One reason for the mixed evidence is that it is hard to precisely estimate the betas of individual stocks, particularly when these betas are time-varying and when the number of return observations is small. Fama and MacBeth (1973) propose to group firms together into portfolios based on their characteristics to reduce measurement error in beta. An important drawback of this approach, however, is that aggregating stocks into portfolios leads to a loss of information if the stocks within a given portfolio are not homogeneous in terms of risk exposure. Another problem is that it is not clear how to model time variation in betas, which is important because Ghysels (1998) shows that misspecifying beta dynamics in a conditional asset pricing model can result in pricing errors that are potentially even larger than those produced by a static model. Although robust, nonparametric approaches can be used to model changes in risk exposures, these methods involve a tradeoff between the timeliness and precision of the beta estimates and do not link risk dynamics to underlying economic theory.

Business cycle fluctuations not only affect the betas of individual stocks but also the risk of the market portfolio itself. For example, during the credit crisis of 2008 annualized market volatility, which is the square root of market variance, shot up from 15% in August to more than 75% two months later. Because a change in market risk affects the risk-return tradeoff, it can be a priced risk factor in the intertemporal CAPM (ICAPM) of Merton (1973). While the CAPM assumes that investors only care about the wealth their portfolio generates one period in the future, the ICAPM recognizes that in reality investors also consider the opportunities they will have for investing this wealth at the end of the period. Thus, in the ICAPM, risk-averse investors want to hedge the exposure of asset returns to changes in state variables that affect future investment opportunities. This hedging demand has an impact on equilibrium returns. Intuitively, investors prefer stocks that pay off when expected future market returns fall or expected future market variances rise. Consequently, they are willing to accept lower returns on these securities, which is the premium they pay to insure against a deterioration of the risk-return tradeoff in the future.

Because portfolio theory shows that the variance of the market portfolio depends on the variances and pairwise correlations of all stocks in the market, any changes in market risk must be driven by changes in these individual variances and correlations. Hence, if market variance risk is priced in the cross-section of stock returns, individual variance risk, correlation risk, or both should be priced. Stocks that covary positively with shocks to average individual variance or with a sudden increase in market-wide correlations should have a positive hedging demand and lower expected returns, which translates into a negative price of variance and correlation risk. Furthermore, since news can have a different impact on asset prices in the long run than in the short run, volatilities and correlations exhibit long-term and short-term movements. For instance, high-frequency movements in risk may reflect liquidity concerns whereas low-frequency components of risk are affected by structural changes in the economic environment and in firm-specific conditions. As a result, long and short run variance and correlation factors may capture orthogonal sources of risk and carry a different price of risk.

1.2 Investor Irrationality in Financial Markets

The systematic deviations of security prices from the predictions of the efficient market hypothesis led to the emergence of an alternative perspective on financial markets in the 1980s, known as behavioral finance.² This view attributes anomalous patterns in asset returns to investor irrationality. The field is based on two building blocks, which are limited arbitrage and behavioral biases. Because of limits to arbitrage it can be risky and costly for rational arbitrageurs to implement strategies to correct mispricing. As a result, they cannot completely eliminate the distorting influence of irrational investors on security prices. The most important risks arbitrageurs face are fundamental risk and noise trader risk. Fundamental

²Shleifer (2000) and Shiller (2005) offer concise introductions to behavioral finance.

risk arises when no perfect substitutes are available for assets that are mispriced, which prevents arbitrageurs from creating a riskless hedge. Noise trader risk is the possibility that mispricing deepens before it is eventually corrected, due to investors who trade on noise instead of real information. When arbitrageurs have limited funds, they can be forced to liquidate their positions before security prices revert to fundamental values. Other factors that impede arbitrage and allow mispricing to persist include transaction costs and short-sales constraints.

The second cornerstone of behavioral finance, behavioral biases, refers to deviations from full rationality investors suffer from. Experimental evidence from psychology shows that biases arise both in the way people form beliefs and in making decisions based on these beliefs. Since the world is a complex place and people's cognitive abilities are limited, they adopt rules of thumb to guide their decision making. Kahneman and Tversky (1974) show that while these heuristics can be very helpful in some situations, they can lead to systematic biases in people's judgements in other cases. For instance, when confronted with new information, people do not update their beliefs as much as they should according to Bayes' rule. Furthermore, their estimates of probabilities can be distorted when they identify patterns that do not exist because an event seems to resemble their image of another event. Other mistakes in beliefs are due to overconfidence, which leads people to overestimate the precision of their judgements. Initial evidence that people do not always act rational when making decisions under uncertainty was presented in an influential paper by Kahneman and Tversky (1979). They set up experiments and find significant deviations from the predictions of expected utility theory. For instance, people exhibit loss aversion, which means that they are more sensitive to losses than to gains. Moreover, they care about changes in wealth relative to a reference point, not about final wealth.

Although these biases only affect security prices when arbitrage is limited and when the strategies of noise traders are correlated, they have a direct impact on the trading behavior and performance of investors. For instance, Barber and Odean (2000) argue that because of overconfidence individual investors trade too much, which leads to poor performance due to high transaction costs. Other studies attribute excessive trading to gambling motives (Kumar (2009b)) and identify important differences in the performance of different groups of investors related to gender and wealth (Barber and Odean (2001)). Poteshman and Serbin (2003) show that people not only exhibit irrational behavior in stock markets but also make errors when trading more complex securities like options.

In the past, neoclassical finance and behavioral finance were usually considered to be mutually exclusive, but recently some academics have started to reconcile both perspectives in new theoretical models. An example is the framework proposed by Barberis, Huang, and Santos (2001), who incorporate findings from psychological experiments, captured by the prospect theory of Kahneman and Tversky (1979), in the traditional consumption-based asset pricing model of Lucas (1978). Others however, like Ross (2005), strongly oppose this view and argue that behavioral finance concepts are not useful because they are not based on a unifying framework like the efficient market paradigm.

1.3 Aim and Outline of the Thesis

The main objective of this dissertation is to shed light on the causes and consequences of time-varying risk in stock markets and to improve the specification and estimation of asset pricing models that incorporate these risk dynamics. Precise measurement of risk is essential for risk managers in financial institutions, for managers in companies to calculate the cost-of-capital used in capital budgeting decisions, for portfolio managers to construct efficient portfolios, and for institutional investors to compute the risk-adjusted performance of their portfolios. Accurate measures of risk are also crucial for research. Important academic applications include the computation of risk-adjusted returns in event studies, the testing of asset pricing models, and the evaluation of the performance of institutional investors like hedge funds, mutual funds, and pension funds.

Another aim of this thesis is to better understand the trading behavior and performance of individual investors, which is important in a world in which people become more responsible for their own financial well-being. The relevance of this research topic is highlighted by the financial crisis in 2008, which showed that many people do not fully understand the structure of complex financial products. This lack of knowledge can have serious consequences for society.

The thesis is organized as follows. Chapter 2 studies the risk dynamics and pricing of 25 stock portfolios in the Fama and French (1993) three-factor model. The portfolios are constructed by sorting stocks from 16 European markets on firm size and book-to-market equity. Time-varying betas are modeled as a linear function of a set of conditioning variables, which includes firm characteristics and macroeconomic variables. The empirical analysis first addresses the question whether factor loadings in the Fama and French (1993) model are time-varying. Subsequently, the performance of conditional specifications of the model in explaining time variation in returns is compared to the performance of their static counterpart. Specifically, for each portfolio time series regressions are run to determine whether pricing errors are zero. Finally, a cross-sectional analysis is performed to identify the sources of any mispricing and to examine whether conditioning can improve the cross-sectional pricing ability of the three-factor model.

In Chapter 3 the analysis of conditional asset pricing models is extended and refined by improving both the specification and estimation of time-varying market betas for a large panel of stocks listed on the NYSE and AMEX. Time variation in CAPM betas is modeled by combining a parametric specification based on economic theory with a nonparametric approach based on data-driven filters. The optimal mix of the two methods is allowed to vary across stocks and over time depending on firm-specific and market-wide conditions. The precision of beta estimates is increased by setting up a hierarchical Bayesian panel model that imposes a common structure on the parameters to exploit information in the cross-section of stocks. Furthermore, high-frequency data is used to estimate realized betas with greater precision and optimal weights are given to past data to increase the timeliness of the estimates. Chapter 3 then investigates whether the time-varying beta estimates from the panel model increase the cross-sectional explanatory power of

the CAPM relative to time-varying betas that are based on traditional specifications and relative to static betas. In addition, the mixed betas from the panel model are used to forecast the covariance matrix of stock returns and to construct minimum variance portfolios, whose out-of-sample performance is compared to that of portfolios formed by existing methods.

Chapter 4 considers the pricing of long and short run variance and correlation risk in aggregate returns and in the cross-section of individual returns. Option data for the S&P 100 index and its constituents is used to calculate implied variances and implied correlations. High-frequency data is used to compute realized variances and correlations, which are then subtracted from their implied counterparts to obtain variance and correlation risk premia. Subsequently, time series regressions are run to determine whether the market variance risk premium, average individual variance risk premium, and correlation risk premium have predictive power for market returns. The second part of chapter 4 focuses on the cross-sectional pricing of long- and short-term volatility and correlation risk by decomposing market volatility, idiosyncratic volatility, and correlations into high- and low-frequency components using a semi-parametric approach. It is tested whether innovations in long and short run components of market volatility, average idiosyncratic volatility, and market-wide correlations carry a significant price of risk in a cross-section of individual stocks. The pricing model also includes size, value, momentum, and liquidity factors to determine whether the variance and correlation factors have explanatory power beyond that of traditional factors.

Chapter 5 examines the impact of option trading on the performance of individual investors. The analysis is performed using a unique database that comprises more than 68,000 accounts and eight million trades in stocks and options at a large online broker in the Netherlands. Investor behavior and performance is studied for the period January 2000 to March 2006, which includes the boom and bust of the tech bubble. The performance of option traders is compared to that of investors who only trade stocks. The analysis accounts for time-varying risk and style tilts in investor portfolios by using dynamic performance evaluation models. The return comparison also controls for known determinants of individual investor performance by attributing returns to a number of investor characteristics. Eventually, return differences between the two groups of investors are related to market timing and trading costs. Chapter 5 then looks into the motivations for trading options. Specifically, option trades are linked to the common stock holdings of investors to determine whether hedging is an important reason for trading options or whether investors use options for gambling because of the leverage they provide and their skewed payoffs. Finally, it is tested whether some option investors possess skills that allow them to consistently outperform others. Subsequently, the characteristics of these successful traders and their portfolios are documented.

Chapter 6 summarizes and discusses the main findings of the dissertation and provides directions for future research. Finally, a note on the structure of the thesis and the notation that is used. Some derivations and proofs are relegated to technical appendices at the end of the corresponding chapter. The notation is consistent within each chapter but can differ slightly across chapters.

Chapter 2

Conditional Asset Pricing and Stock Market Anomalies in Europe*

This chapter provides European evidence on the ability of static and dynamic specifications of the Fama and French (1993) three-factor model to price 25 size and book-to-market sorted portfolios. In contrast to U.S. evidence, we detect a small-growth premium and find that the size effect is still present in Europe. Furthermore, we document strong time variation in risk factor loadings. Incorporating these risk fluctuations in conditional specifications of the three-factor model clearly improves its ability to explain time variation in expected returns. However, the model still fails to completely capture the cross-sectional variation in stock returns as it is unable to explain the momentum effect.

2.1 Introduction

The Capital Asset Pricing Model (CAPM) has been one of the cornerstones of modern finance since its development in the 1960s. However, starting in the eighties a number of patterns in the cross-section of average returns have been detected that question the validity of the model, including the size effect (Banz (1981)), value premium (Basu (1977)), and momentum effect (Jegadeesh and Titman (1993)). In addition, Fama and French (1992) show that size and book-to-market equity (B/M) are better able to capture the cross-section of returns than market beta. While the Fama and French (1993) three-factor model seems to be able to repair most of the cracks in the building of modern finance, Fama and French (1996) show that it is unable to explain the momentum effect. Apart from these cross-sectional

*This chapter is based on Bauer, Cosemans, and Schotman (2009), forthcoming in *European Financial Management*.

anomalies, prior research has also found that firm characteristics (Lewellen (1999)) and macroeconomic variables (Ferson and Harvey (1999)) predict significant time variation in expected returns on size-B/M sorted portfolios.

Rational asset pricing theory posits that these variables have predictive power because they capture information about time-varying risk. Accordingly, static models fail to explain the cross-section of returns because they ignore differences in risk dynamics across stocks. Therefore, much recent work has focused on *conditional* asset pricing models, in which risk loadings are allowed to vary over time. Since the empirical results of these studies are mixed and primarily based on U.S. data, our main goal in this chapter is to provide out-of-sample evidence on the performance of the conditional three-factor model. Specifically, we test whether factor loadings are time-varying and if so, to what extent dynamic specifications of the model explain time variation and cross-sectional variation in returns on 25 size-B/M portfolios constructed using stocks from 16 European markets.

We first examine the predictive power of a set of macroeconomic and portfolio-specific variables for European size-B/M portfolios. Next, we test whether portfolio betas are time-varying by modeling variation in risk loadings as a function of the predictive variables. Subsequently, we investigate whether conditional models completely explain conditional expected returns, i.e., whether conditional alphas are zero. We also test the weaker hypothesis that conditional alphas are constant. Finally, we combine the time series analysis with the cross-sectional framework developed by Brennan, Chordia, and Subrahmanyam (1998). Specifically, we calculate risk-adjusted returns and perform a cross-sectional analysis to examine whether cross-sectional variation in pricing errors is related to size, book-to-market, and past returns.

We extend existing work in several ways. First, by using a large data set of European stocks we provide out-of-sample empirical evidence on the time-varying behavior of risk and the performance of conditional asset pricing models. Following Fama and French (2006), we construct the 25 size-B/M portfolios and the risk factors using merged data from the markets of interest, which enables us to form well-diversified portfolios. Thus, we adopt a pan-European approach motivated by the increasing integration of European markets that started in the mid-1980s (Bekaert, Hodrick, and Zhang (2009a) and Eiling and Gerard (2007)), which coincides with the beginning of our sample. In contrast, Fama and French (1998) examine the presence of a value premium in international markets by forming portfolios for each country separately. An important drawback of this method is that for many countries the resulting portfolios only contain a few stocks and, consequently, that idiosyncratic risk is not fully eliminated.

We find that for our sample period the explanatory power of the three-factor model in Europe is higher than in the United States, providing further support for the pan-European perspective we take. Rouwenhorst (1998) also takes a pan-European perspective by aggregating stocks into momentum portfolios. He confirms that the profitability of momentum strategies is not driven by country effects. Heston, Rouwenhorst, and Wessels (1999) document a negative relation between European stock returns and firm size.

Our work contributes to these studies as follows. First, we allow for time-varying risk in a conditional three-factor model whereas Fama and French (1998), Heston, Rouwenhorst, and Wessels (1999), and Rouwenhorst (1998) control for risk using unconditional two-factor models. Second, we focus on multiple anomalies simultaneously while they only consider a single anomaly. Third, we examine a more recent time period, which includes the burst of the tech bubble.

On the methodology side, the time series tests we employ avoid the problems associated with cross-sectional tests discussed by Lewellen, Nagel, and Shanken (2009), because they impose the theoretical restrictions that risk premia should equal expected excess factor returns and that the zero-beta rate should be equal to the risk-free rate.

Our empirical results show important differences between U.S. and European data, which motivates the analysis of the performance of conditional asset pricing models in Europe. In particular, we find that the size effect, which has vanished in the United States after its discovery (Cochrane (2005)), is still present in Europe. In addition, our time series analysis reveals that the unconditional three-factor model leaves significant pricing errors. Strikingly, the small-growth portfolio, known to be hard to price in the United States because it generates significantly negative alphas, produces significantly positive pricing errors in Europe. We also find that macroeconomic and portfolio-specific variables have substantial predictive power for returns on the size-B/M portfolios. Using these variables as instruments for conditional betas, we document strong evidence of time-varying risk. Incorporating these fluctuations in risk improves the performance of the three-factor model in explaining time variation in portfolio returns. Nevertheless, even after allowing for variation in factor loadings, the three-factor model does not completely explain conditional expected returns as pricing errors for some portfolios remain predictable.

Our cross-sectional findings show that the rejection of the three-factor model is due to strong momentum effects in returns on the 25 size-B/M portfolios. Both the static and dynamic three-factor model do not capture the explanatory power of past return (momentum) variables for the cross-section of portfolio returns. Our European evidence supports the U.S. findings of Ferson and Harvey (1999), Avramov and Chordia (2006a), Lewellen and Nagel (2006), and Petkova and Zhang (2005). In particular, although betas do vary over time, these fluctuations are too small to explain the momentum effect.

The remainder of this chapter is organized as follows. Section 2.2 provides the rationale for conditional asset pricing models and reviews existing empirical evidence. Section 2.3 explains our methodology and Section 2.4 describes the data set. In Section 2.5 we present our empirical findings and discuss the robustness of the results. Section 2.6 concludes.

2.2 Conditional Asset Pricing: Theory and Evidence

Proponents of conditional asset pricing theory argue that the failure of unconditional models to explain the cross-section of stock returns might be due to their assumption that risk loadings remain constant over time. Santos and Veronesi (2004) show within a general equilibrium model that market betas vary substantially when the covariation between a firm's cash flows and the aggregate economy is large. If true betas are time-varying, static models are misspecified and give an incomplete description of stock returns. Indeed, abundant empirical evidence of time variation in beta has been found, which has motivated the development of conditional models that allow factor loadings to vary. Hansen and Richard (1987) provide theoretical support for dynamic models by showing that a conditional CAPM can hold perfectly even if its unconditional counterpart fails.

According to conditional asset pricing theory, a significant relation between predictive variables and the time series and cross-section of returns must be due to their association with risk. In particular, the variables must contain information about time variation in risk and consequently in expected returns. Furthermore, differences in risk dynamics across stocks induce cross-sectional variation in conditional expected returns. This implies that the power of the predictors should disappear once we adequately control for fluctuations in risk. Gomes, Kogan, and Zhang (2003) show theoretically that the ability of size and B/M to explain cross-sectional variation in returns is due to their correlation with the true conditional market beta. Zhang (2005) extends this work and argues that because of costly reversibility of capital, value firms have countercyclical betas while betas of growth stocks are procyclical. Because the price of risk is also countercyclical his model can explain the value premium within a rational framework.

In contrast, the mispricing view put forward by Lakonishok, Shleifer, and Vishny (1994) and Daniel and Titman (1997) asserts that the significant association between predictive variables and expected returns is related to investor cognitive biases. Specifically, this story says that the predictors contain information about mispricing of securities and, consequently, that their predictive power will persist even when risk fluctuations are taken into account.

Hitherto, empirical evidence on the performance of conditional asset pricing models is inconclusive and primarily based on U.S. data. Favorable results are documented by Jagannathan and Wang (1996), who find that a conditional CAPM extended by a proxy for the return on human capital leaves insignificant pricing errors when applied to portfolios sorted on size and beta. Lewellen (1999) shows that after controlling for its role as determinant of conditional betas, B/M contains little incremental information about time variation in expected returns. Ferson and Harvey (1998) find that the cross-sectional explanatory power of firm-specific attributes like book-to-market mainly arises from their role as instruments for risk rather than from their relation to mispricing. Wu (2002) argues that a conditional three-factor model captures momentum and reversal effects in U.S. stock returns

when the assumption of a linear relation between conditioning variables and betas is relaxed.¹

More recently, Ang and Chen (2007) document strong evidence of time variation in betas of portfolios sorted on B/M. They find that a conditional CAPM in which time-varying betas are treated as latent state variables is able to capture the book-to-market effect. Adrian and Franzoni (2009) propose a conditional CAPM in which investors learn about unobserved time-varying risk by observing realizations of returns. Their learning CAPM cannot be rejected when applied to size-B/M portfolios.

In contrast, results found by other studies are less favorable. Ferson and Harvey (1999) show that even in a conditional three-factor model, proxies for time variation in expected returns based on macroeconomic instruments have significant cross-sectional explanatory power for returns on size-B/M sorted portfolios. Petkova and Zhang (2005) confirm empirically the theoretical prediction of Zhang (2005) that value firms are riskier than growth firms in economic downturns when the expected market premium is high. However, they note that the covariance between beta and the market risk premium is too small to explain the magnitude of the value premium. Lewellen and Nagel (2006) also find that this covariance is insufficient to explain the large unconditional alphas produced by book-to-market and momentum portfolios.

Using individual stocks as test assets, Avramov and Chordia (2006a) find that conditional multifactor models can explain size and value anomalies but are unable to capture momentum and turnover effects in returns. Lewellen, Nagel, and Shanken (2009) argue that the favorable evidence on the ability of conditional models to price size-B/M sorted portfolios is largely due to the low power of the cross-sectional tests employed in many papers. In particular, they show that when risk premia are unrestricted, *any* factor that is only weakly correlated with *SMB* and *HML* will price the size-B/M portfolios due to their strong factor structure.

2.3 Methodology

2.3.1 Time Series Test Methodology

The conditional three-factor model can be written as

$$E_t(R_{it+1}) = \alpha_{it} + \sum_{k=1}^3 \beta_{ikt} E_t(FF_{kt+1}), \quad (2.1)$$

where R_i is the excess return on asset i , FF is a vector containing the three Fama-French factors R_M , SMB , and HML , $E_t(\cdot)$ is the conditional expectation, given

¹Wu (2002) tests the conditional linear-exposure three-factor model using standard OLS time series regressions but employs a cross-sectional GMM test for the conditional nonlinear-exposure three-factor model. This makes it difficult to determine whether his results are driven by the difference in model specification or by a difference in the power of both tests.

the public information set at time t , and β_{ikt} is the conditional beta with respect to the k 'th factor.

Following Shanken (1990), we model time variation in alphas and betas by allowing them to depend linearly on a set of predetermined instruments (conditioning variables).² This approach explicitly links conditional betas to observable state variables, consistent with the economic motivation for conditional models. In particular, in this framework conditional betas are given by

$$\beta_{ikt} = \gamma_{ik0} + \gamma_{ik1}Z_{it}, \quad (2.2)$$

where γ_{ik0} is a scalar, γ_{ik1} a vector of L parameters, and Z_{it} a vector of L instruments. We test the hypothesis that risk loadings are constant over time by examining whether the γ_{ik1} parameters are equal to zero. Analogous to the specification of conditional betas, the conditional alpha is given by

$$\alpha_{it} = \alpha_{i0} + \alpha_{i1}W_{it}, \quad (2.3)$$

where W_{it} is a vector of instruments for alpha.

We test the hypothesis that the three-factor model completely explains conditional expected portfolio returns, which corresponds to the null hypothesis that the conditional alpha in equation (2.1) is equal to zero. Thus, rational asset pricing theory predicts that the α_{i0} and α_{i1} parameters in equation (2.3) should all be zero. We also test the weaker condition that alpha is constant over time, i.e., that the α_{i1} are zero. Under this null hypothesis, the instruments do not predict expected portfolio returns after their role as instrumental variables for conditional risk loadings is taken into account. The alternative hypothesis is that the conditioning variables are related to time-varying mispricing.

Combining equations (2.1), (2.2), and (2.3) leads to the econometric model

$$R_{it+1} = \alpha_{i0} + \alpha_{i1}W_{it} + (\gamma_{ik0} + \gamma_{ik1}Z_{it})FF_{kt+1} + \epsilon_{it+1}. \quad (2.4)$$

We evaluate alternative model specifications based on (2.4), with both constant and time-varying alphas and betas and various combinations of instrumental variables, using the adjusted R^2 and Akaike information criterion.

2.3.2 Cross-Sectional Framework

In the cross-sectional tests we examine the predictive power of various non-risk characteristics. Rational asset pricing theory predicts that non-risk security characteristics like size and B/M should not have any cross-sectional explanatory power for returns incremental to the risk factors included in the asset pricing model. This

²Other approaches to allow for time variation in betas include estimating conditional betas using short-window regressions (Lewellen and Nagel (2006)), rolling regressions (Fama and French (1997)), and modeling beta as a latent autoregressive process (Ang and Chen (2007) and Ammann and Verhoben (2008)). In contrast to the approach we use, these methods do not explicitly link risk dynamics to business cycle variations.

hypothesis can be tested by running the following cross-sectional regression,

$$R_{it+1} = c_{0t+1} + \sum_{k=1}^3 \lambda_{kt+1} \hat{\beta}_{ikt} + \sum_{m=1}^M c_{mt+1} P_{mit} + \nu_{it+1}, \quad (2.5)$$

where λ_k is the risk premium for factor k , c_m is the reward to non-risk characteristic m , and P_{mi} is the value of characteristic m for portfolio i . The null hypothesis that expected returns on portfolio i only depend on its sensitivity to the risk factors in the model, measured by $\hat{\beta}_{ik}$, implies that all loadings c_m on the non-risk factors must be zero. As noted by Berglund and Knif (1999), the traditional Fama and MacBeth (1973) two-step procedure commonly used to test this hypothesis suffers from an errors-in-variables problem, since the betas included as regressors in the second-stage cross-sectional regressions are estimated with error in the first-stage time series regressions.

In order to circumvent this problem, we follow Brennan, Chordia, and Subrahmanyam (1998) and Avramov and Chordia (2006a) by regressing *risk-adjusted* returns obtained from the time series regression (2.4) on the portfolio characteristics size, book-to-market, and cumulative past returns. Specifically, the estimated risk-adjusted return is given by

$$R_{it+1}^* \equiv R_{it+1} - \sum_{k=1}^3 \hat{\beta}_{ikt} FF_{kt+1}. \quad (2.6)$$

We calculate the risk-adjusted return R_{it+1}^* in (2.6) as the sum of the pricing error α_{it} and the error term ϵ_{it+1} from the first-pass time series regression (2.4). This risk-adjusted return is then used as dependent variable in the second-stage cross-sectional regression,

$$R_{it+1}^* = c_{0t+1} + \sum_{m=1}^M c_{mt+1} P_{mit} + \eta_{it+1}. \quad (2.7)$$

This approach avoids any errors-in-variables bias because the betas estimated in the time series regressions do not show up as explanatory variables in the cross-sectional regressions. The relation between the time series tests and the cross-sectional analysis is as follows: if the time series regressions identify significant pricing errors, the cross-sectional tests reveal whether this mispricing is related to size, value, or momentum effects.

We test the hypothesis that expected returns only depend on the risk characteristics of returns by calculating the Fama-MacBeth (FM) estimator for the non-risk characteristics, which is the time series average of the monthly parameter estimates c_{mt} . The standard error of the FM estimator is also calculated from the time series of these monthly estimates.

2.4 Data and Descriptive Statistics

The MSCI data set we use consists of the monthly return and book and market value for a sample of common stocks from 16 European countries that covers approximately 80% of European stock market capitalization. All variables are denominated in Euros.³ The raw data set includes 2503 firms and covers the period from February 1985 to June 2002. The stocks are listed on the exchanges of Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, the Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and the UK. The number of stocks per country ranges from 37 for Ireland to 519 for the UK. The MSCI data is free from survivorship bias as it includes historical data for firms that are delisted over time. Furthermore, historical data for newly included stocks is not added to the data set, which should prevent any backfilling bias.

A stock is used in our analysis in a given month t if it satisfies the following criteria: (i) data should be available in month $t - 1$ for size as measured by market capitalization and for the book-to-market ratio. This condition is imposed because returns for month t are calculated for portfolios formed at the end of month $t - 1$ on size and book-to-market; (ii) its book-to-market equity is non-negative. This last requirement follows Fama and French (1993). The screening process leads to a sample that contains on average 1315 stocks per month. The total number of stocks over the full sample period is 2165. Since we need portfolio returns over the past 12 months to calculate cumulative lagged returns as a proxy for momentum, the analysis starts in February 1986 and ends in June 2002.

Our test assets are 25 portfolios formed on size and B/M, which have become standard in asset pricing tests after the failure of the CAPM to explain size and B/M effects in returns. Following Fama and French (2006), we use merged data from all countries to construct the portfolios. Table 2.1 reports the size and B/M characteristics of the 25 European portfolios. We also show the characteristics of the 25 U.S. size-B/M portfolios for the same period.⁴ To facilitate the comparison, we express the market capitalization of the European portfolios in U.S. dollars. Although these statistics show that the firms in the smallest quintile of the European portfolios are somewhat larger than those in the corresponding U.S. quintile, the general size and B/M characteristics of the European sample closely resemble those of the U.S. portfolios. Thus, the relative interpretation of small and large firms and of value and growth stocks is similar in both samples.⁵

The upper left corner of Table 2.2 presents average monthly returns for the 25

³Converting local currency returns to returns in one base currency follows previous international asset pricing studies (e.g., Fama and French (1998) and Rouwenhorst (1998)). In addition, in the robustness section we eliminate any country tilts in the portfolios. As a result, even if exchange rate risk plays a role, it affects each portfolio in a similar way and therefore any effect on our empirical results will be small.

⁴Data for the 25 U.S. size-B/M portfolios is obtained from Kenneth French's data library: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

⁵In the robustness section we purge the size and B/M characteristics of each firm from country effects before constructing the portfolios. This ensures that also *within* Europe firm size and B/M are comparable across countries.

Table 2.1: **Characteristics of 25 Portfolios Formed on Size and B/M: Europe and the United States**

This table presents average size and book-to-market characteristics of 25 size-B/M portfolios for the period February 1986 through June 2002. The portfolios are constructed by first sorting stocks independently into size and B/M quintiles. The 25 size-B/M portfolios are then formed as the intersections of the five size and the five B/M quintiles. The left hand side of the table reports average market capitalization in billions of U.S. dollars and the average B/M ratio for the 25 European portfolios. The right hand side of the table presents average market capitalization in billions of U.S. dollars and the average B/M ratio for the 25 U.S. portfolios.

Size Quintiles	Book-to-Market Equity (B/M) Quintiles				
	Low	2	3	4	High
	Europe				
Small	Average Market Cap (\$ Billions)				
	0.11	0.11	0.11	0.11	0.09
	0.29	0.29	0.30	0.29	0.28
	0.61	0.61	0.62	0.60	0.60
	1.50	1.46	1.48	1.46	1.42
Big	17.01	10.97	8.02	6.50	6.33
	Average B/M Ratio				
Small	0.16	0.35	0.54	0.79	2.03
	0.17	0.35	0.53	0.78	1.71
	0.17	0.35	0.53	0.77	1.53
	0.17	0.35	0.53	0.77	1.40
	0.09	0.35	0.53	0.77	1.40
	Average B/M Ratio				
	0.21	0.43	0.59	0.77	1.40
	0.20	0.40	0.57	0.75	1.22
	0.20	0.39	0.56	0.77	1.19
	0.20	0.40	0.57	0.78	1.17
	Average B/M Ratio				
	0.19	0.39	0.56	0.77	1.11

European size-B/M portfolios. For comparison, in the upper right corner we also report statistics for the 25 U.S. size-B/M portfolios for the same period. Strikingly, in the European sample the small-growth portfolio (S1/B1) has the highest average return. In contrast, Fama and French (1996) find that in the United States the return on the small-growth portfolio is the lowest of all 25 portfolios, which is confirmed by the statistics we report for the U.S. portfolios. Table 2.2 also shows the presence of a size effect in the European sample, which is absent in U.S. data. The value premium is positive in both samples but insignificant. Overall, Table 2.2 reveals important differences between European and U.S. data, which motivates our analysis of the performance of conditional asset pricing models in Europe.

An important issue when applying asset pricing models to European stock markets is whether country-specific or pan-European versions of the models should be used. Bekaert, Hodrick, and Zhang (2009a) and Eiling and Gerard (2007) find evidence of integration of European capital markets from the mid-1980s onwards, which coincides with the beginning of our sample period. This motivates the construction of the Fama-French risk factors on a pan-European level, consistent with our portfolio formation procedure. We choose the MSCI Europe index as a proxy for the market portfolio because of its broad coverage of European stock market capitalization. We subtract the three-month German FIBOR rate as a proxy for the risk-free rate to obtain the market premium R_M . We follow the procedure outlined by Fama and French (1993) for constructing the *SMB* and *HML* factors on a European level.

We use both macroeconomic and portfolio-specific variables as instruments for conditional alphas and betas because of their predictive power for returns. Macroeconomic variables shown to predict returns include the default spread (*DEF*; Keim and Stambaugh (1986)), the risk-free rate (R_F ; Fama and Schwert (1977)), and the term spread (*TERM*; Fama and French (1989)).⁶ We define the default spread as the yield spread between Moody's Baa- and Aaa-rated corporate bonds. The term spread is the difference between the yield on ten-year German government bonds and the three-month FIBOR rate. The theoretical motivation for choosing the portfolio-specific variables size and book-to-market as instruments is given by Gomes, Kogan, and Zhang (2003), who show that the ability of size and B/M to explain the cross-section of returns is due to their correlation with the true conditional market beta. In particular, they demonstrate that size captures the component of a firm's systematic risk related to its growth options whereas the book-to-market ratio is a measure of the risk of the firm's assets in place.

Jagannathan and Wang (1996) stress that, although several variables may help predict the business cycle, we have to restrict ourselves to a small number of such variables to ensure precision in the estimation of the model parameters. We use the adjusted R^2 and Akaike information criterion to determine the optimal set of conditioning variables. These model selection criteria prefer the default spread,

⁶Because of their predictive power for market returns, innovations in these macroeconomic variables could be priced risk factors in the intertemporal CAPM of Merton (1973). However, due to measurement errors in economic data, these macroeconomic factors have lower explanatory power for the cross-section of returns than the Fama and French (1993) factor-mimicking portfolios.

Table 2.2: Average Returns on European, U.S., and Country- or Sector-Neutral European Size-B/M Portfolios

This table presents value-weighted average monthly returns on 25 size-B/M stock portfolios for the period February 1986 through June 2002. The portfolios in the upper left corner are the original European portfolios constructed by sorting stocks each month independently into size and B/M quintiles. The 25 size-B/M portfolios are formed as the intersections of the five size and the five B/M quintiles. The portfolios in the upper right corner are the 25 U.S. size-B/M portfolios for the same period. The bottom left and bottom right corners show returns for European size-B/M portfolios constructed using country-neutral and sector-neutral firm size and book-to-market, respectively. Value-weighted returns for month $t + 1$ are calculated for portfolios formed at the end of month t as the value-weighted average of the excess returns on the individual stocks in the portfolios. $H - L$ is the value premium for a given size quintile, defined as the average of the time series of monthly differences between the return for the highest B/M quintile and the return for the lowest B/M quintile within a size group. Similarly, $S - B$ is the size premium for a given B/M quintile, defined as the average of the time series of monthly differences between the return for the smallest size quintile and the return for the largest size quintile within a B/M group. $t(H-L)$ and $t(S-B)$ are the average monthly differences divided by their standard error.

Size Quintiles	Book-to-Market Equity (B/M) Quintiles													
	Low	2	3	4	High	H-L	t(H-L)	Low	2	3	4	High	H-L	t(H-L)
Europe														
Small	2.35	1.19	0.60	0.92	1.20	-1.15	-2.19	-0.24	0.72	0.86	1.11	0.98	1.22	3.34
2	1.00	0.70	0.54	0.87	1.19	0.19	0.48	0.22	0.60	0.87	0.92	0.84	0.62	1.84
3	0.48	0.44	0.36	0.79	1.16	0.68	2.06	0.39	0.69	0.70	0.80	1.03	0.64	1.57
4	0.41	0.34	0.56	0.87	1.08	0.67	2.04	0.74	0.75	0.75	0.90	0.85	0.11	0.30
Big	0.52	0.45	0.55	0.72	0.82	0.30	0.82	0.76	0.79	0.71	0.73	0.65	-0.10	-0.34
S-B	1.83	0.73	0.05	0.20	0.38			-0.99	-0.06	0.15	0.39	0.33		
t(S-B)	3.82	2.34	0.17	0.86	1.17			-2.08	-0.13	0.40	1.05	0.91		
Europe: Country-Neutral														
Small	2.18	0.81	0.80	0.76	0.98	-1.20	-2.58	2.26	1.21	1.02	0.90	1.17	-1.09	-2.32
2	0.83	0.47	0.44	0.65	0.97	0.14	0.40	1.08	0.66	0.87	0.89	1.15	0.07	0.20
3	0.51	0.49	0.27	0.70	0.91	0.40	1.29	0.34	0.54	0.61	0.78	1.19	0.84	2.61
4	0.43	0.44	0.57	0.81	0.76	0.34	0.99	0.35	0.51	0.51	0.83	0.84	0.49	1.63
Big	0.56	0.51	0.58	0.71	0.57	0.02	0.06	0.45	0.50	0.76	0.57	0.45	0.00	0.01
S-B	1.63	0.30	0.22	0.05	0.40			1.81	0.71	0.26	0.33	0.72		
t(S-B)	3.82	1.06	0.88	0.20	1.35			4.29	2.36	1.06	1.30	2.55		

size, B/M, and interaction terms between default spread and size and between default spread and B/M as instruments for conditional betas. The interaction terms allow the relation between risk and portfolio characteristics to vary over time as a function of the business cycle, consistent with the economic motivation for conditional asset pricing models. For modeling conditional alphas, the information criteria indicate that including the default spread, risk-free rate, term spread, and the two portfolio-specific instruments size and B/M leads to the best model fit. Apart from statistical considerations, choosing a different set of conditioning variables for alpha and beta is also in line with Avramov and Chordia (2006a).

For every portfolio the following non-risk characteristics are calculated each month as possible determinants of the cross-section of expected returns:

- *SIZE*: average market value of equity of the firms in the portfolio in billions of euros, included to assess the significance of the size effect.
- *B/M*: sum of book equity for the firms in the portfolio divided by the sum of their market capitalization, included to examine the significance of the value premium.
- *RET2-3*, *RET4-6*, *RET7-12*: cumulative portfolio returns over the second through third, fourth through sixth, and seventh through twelfth months prior to the current month, respectively, included to analyze the significance of the momentum effect.

Because the distributions of these characteristics display considerable skewness, we use their logarithmic transformations. Following Brennan, Chordia, and Subrahmanyam (1998), in the cross-sectional analysis we normalize the characteristics by expressing them as deviations from their cross-sectional means. Thus, for the *average* portfolio the value of each characteristic is zero. Consequently, under both the null hypothesis that non-risk characteristics do not have significant incremental power for capturing the cross-section of returns and the alternative hypothesis that they do have significant explanatory power, the return on the *average* portfolio only depends on its risk characteristics.

2.5 Empirical Results

2.5.1 Time Series Evidence on the Unconditional Three-Factor Model

As a benchmark we first consider time series regression results for the unconditional three-factor model, shown in Table 2.3. The explanatory power of the model in terms of adjusted R^2 is high, ranging from 73% to 89%. In fact, for our sample period the three-factor model has higher explanatory power for size-B/M portfolios in Europe than in the United States. This result lends support to the assumption that European financial markets are fairly integrated and justifies our use of a pan-European version of the model.

The intercept in the time series regressions is the pricing error, which represents the portion of the excess portfolio return left unexplained by the risk factors in the model. If the three-factor model completely explains the cross-section of average returns, the intercepts in the time series regressions should all be zero. The empirical results show that the intercept is significant at the 5% level for four out of 25 portfolios.

Notably, the pricing error of the small-growth portfolio is significantly *positive* and large in economic terms, whereas in the United States this portfolio produces significant *negative* pricing errors (Fama and French (1996)). Thus, the small-growth anomaly we observe in Europe is exactly opposite to that in the United States. Moreover, Fama and French (1998) document a value premium in international data. However, their sample period (1974 - 1994) excludes the tech bubble in the late nineties. A subperiod analysis confirms that the alphas for the small-growth portfolios S1/B1 and S1/B2 are higher in the second half of our sample period (1994 - 2002), which is not in the sample of Fama and French (1998).⁷ This is illustrated in Figure 2.1, which displays the evolution through time of alphas in the conditional three-factor model for the five portfolios in the smallest quintile. As noted by Cochrane (2005), in the United States the value premium has also decreased substantially during the 1990s. Consistent with these explanations, a closer look at the composition of the S1/B1 portfolio shows that it includes many start-up companies, which are relatively small and have low book-to-market ratios and whose stock prices soared during the tech bubble. Another striking feature of this portfolio is the high first-order autocorrelation of its monthly return ($AR(1) = 0.41$), which could be due to thin trading. Therefore, part of the high return on this portfolio might also be a compensation for liquidity risk.

In general, several pricing errors are quite large in absolute value, particularly for some of the small portfolios. This is confirmed by the Gibbons, Ross, and Shanken (1989) test, which strongly rejects the hypothesis that the alphas of the 25 size-B/M portfolios are jointly equal to zero. These results motivate the extension of the model to conditional specifications.

2.5.2 Predictability of Size and Book-to-Market Portfolio Returns

Before turning to the analysis of the conditional three-factor model, we first examine the time series relation between expected returns on the 25 size-B/M portfolios and the predictive variables R_F , DEF , $TERM$, $SIZE$, and B/M . Table 2.4 summarizes time series regressions of portfolio returns on this set of lagged variables. The results confirm that the variables are significant predictors of time variation in expected portfolio returns. The explanatory power of the predictive variables in terms of adjusted R^2 ranges from 2% for some of the large cap portfolios to 11% for some small cap portfolios, consistent with results documented for the United States by Ferson and Harvey (1999). Moreover, the coefficients on the predictors

⁷In particular, for the subperiod from February 1986 to June 1994 the alpha of portfolio S1/B1 is 1.05, whereas for the subperiod from July 1994 to June 2002 its alpha equals 1.72.

Table 2.3: **Unconditional Three-Factor Model Regressions**

This table reports parameter estimates for unconditional least squares three-factor regressions

$$R_{it} = \alpha_i + \beta_i R_{Mt} + \delta_i SMB_t + \phi_i HML_t + \epsilon_{it}.$$

Monthly excess returns on 25 size-B/M portfolios are regressed on a constant, the market premium R_M , and the factor-mimicking portfolios SMB and HML . $RMSE$ is the root mean squared pricing error. $GRS F$ is the Gibbons, Ross, and Shanken (1989) test statistic for the null hypothesis that the intercepts in the regressions for all portfolios are jointly equal to zero.

	α	β	δ	ϕ	$t(\alpha)$	$t(\beta)$	$t(\delta)$	$t(\phi)$	Adj. R^2
S1/B1	1.47	1.24	1.73	-1.03	4.51	18.21	12.42	-10.92	74.0
S1/B2	0.38	1.04	1.05	-0.20	1.62	21.11	10.42	-2.93	72.6
S1/B3	-0.19	0.92	0.90	0.15	-0.94	21.68	10.38	2.50	73.3
S1/B4	0.10	0.93	0.81	0.32	0.60	26.50	11.30	6.57	80.5
S1/B5	0.15	1.09	1.06	0.59	0.80	28.25	13.41	10.97	83.9
S2/B1	0.25	1.10	1.22	-0.69	1.34	28.25	15.37	-12.72	85.1
S2/B2	-0.07	0.97	0.94	-0.09	-0.45	30.52	14.54	-2.11	84.4
S2/B3	-0.24	0.91	0.81	0.21	-1.68	30.85	13.53	5.08	84.6
S2/B4	0.05	0.90	0.77	0.39	0.38	31.95	13.40	9.93	86.1
S2/B5	0.24	0.99	0.78	0.69	1.53	30.31	11.72	15.19	86.0
S3/B1	-0.21	1.03	0.92	-0.44	-1.34	30.83	13.48	-9.55	85.5
S3/B2	-0.30	0.95	0.68	0.13	-2.38	36.19	12.79	3.45	87.7
S3/B3	-0.40	0.95	0.64	0.23	-2.88	32.82	10.84	5.64	85.5
S3/B4	-0.02	0.99	0.57	0.40	-0.12	33.21	9.28	9.60	86.1
S3/B5	0.20	1.00	0.83	0.68	1.17	28.14	11.47	13.82	84.2
S4/B1	-0.23	1.07	0.66	-0.45	-1.55	34.93	10.61	-10.59	87.7
S4/B2	-0.30	0.93	0.38	0.08	-2.47	37.36	7.39	2.23	87.8
S4/B3	-0.13	0.94	0.40	0.25	-0.93	32.35	6.75	6.32	84.7
S4/B4	0.12	0.99	0.33	0.44	0.82	32.84	5.40	10.62	85.6
S4/B5	0.23	1.04	0.43	0.60	1.29	28.27	5.71	11.77	82.5
S5/B1	0.21	0.99	-0.28	-0.51	1.29	29.78	-4.18	-11.06	84.9
S5/B2	-0.01	0.95	-0.13	-0.03	-0.11	38.97	-2.72	-0.92	88.9
S5/B3	0.01	1.00	-0.16	0.18	0.08	34.32	-2.77	4.52	86.2
S5/B4	0.13	1.00	-0.11	0.33	0.90	33.96	-1.89	8.01	86.1
S5/B5	0.14	1.08	-0.29	0.64	0.62	22.74	-3.03	9.59	75.7
RMSE	0.36								
GRS F	2.47								
p -Value	0.00								

Figure 2.1: Conditional Alphas for Portfolios in Smallest Size Quintile

This figure plots the evolution through time of monthly conditional alphas for the five portfolios in the smallest size quintile. The instruments in the specification of the conditional alpha are the risk-free rate, default spread, term spread, portfolio size, and portfolio B/M. The specification for the conditional risk loadings includes the default spread, portfolio size and book-to-market, and interaction terms between default spread and size and between default spread and B/M. S1/B1 is the portfolio with smallest size and lowest book-to-market ratio and S1/B5 denotes the portfolio with smallest size and highest book-to-market ratio.

vary considerably across portfolios, which suggests that they also have explanatory power for the cross-section of portfolio returns. Table 2.4 further reports regression results for the three Fama-French factors. The predictive power for R_M and SMB is in line with that for the 25 portfolios. However, HML seems not very predictable, suggesting that it contributes little to explaining time-varying expected returns.

Table 2.4: **Predictability of Size and Book-to-Market Portfolio Returns**

Excess portfolio returns and factor premia are regressed on a set of predictive variables. The predictors include the risk-free rate R_F , default spread DEF , term spread $TERM$, portfolio market capitalization $SIZE$, and portfolio book-to-market B/M . $SIZE$ and B/M are expressed as natural logarithms. In the regressions of the three risk factors, $SIZE$ and B/M are the cross-sectional sums of portfolio market capitalization and book-to-market. The sample period is February 1986 to June 2002 and the number of observations is 197. The table reports OLS estimates of the coefficients. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively. The explanatory power of the predictive variables is measured by the adjusted R^2 .

	R_F	DEF	TERM	SIZE	B/M	Adj. R^2
S1/B1	-15.67*	-38.71	-7.89	-19.87***	-14.76***	11.4
S1/B2	-19.29***	-45.69**	-15.30**	-14.25***	-9.59***	10.6
S1/B3	-21.62***	-36.50**	-15.27**	-10.70***	-5.44	11.3
S1/B4	-19.86***	-24.03	-18.40***	-9.79***	-4.19	10.3
S1/B5	-18.21***	-24.33	-13.90*	-10.84***	-2.70	6.5
S2/B1	-14.25**	-64.79***	-7.25	-19.35***	-8.21***	10.8
S2/B2	-13.08***	-27.76	-10.81	-11.67***	-5.58	7.0
S2/B3	-14.32***	-12.41	-13.27**	-9.25***	-2.67	7.5
S2/B4	-15.61***	-2.65	-13.26**	-7.27***	-1.47	6.5
S2/B5	-25.45***	-22.71	-26.84***	-12.37***	-2.70	11.5
S3/B1	-16.02***	-35.42*	-12.28	-12.54***	-3.70	9.1
S3/B2	-18.12***	-17.11	-17.54***	-8.37***	-3.13	6.4
S3/B3	-17.20***	-7.84	-17.37***	-4.97***	1.24	4.3
S3/B4	-21.22***	-3.73	-20.85***	-6.38***	2.71	7.9
S3/B5	-25.04***	-17.45	-25.69***	-6.97***	1.50	8.2
S4/B1	-22.15***	-58.09***	-20.18***	-7.27***	-0.63	7.1
S4/B2	-15.19***	-23.94	-15.12**	-3.53**	2.42	3.6
S4/B3	-15.06***	-21.75	-13.23*	-2.89*	4.02	3.1
S4/B4	-21.25***	-24.26	-21.01***	-4.70***	4.74*	6.4
S4/B5	-29.19***	-41.34*	-26.76***	-6.70***	5.73**	10.7
S5/B1	-12.12**	-5.45	-16.86*	-7.97***	-0.08	4.7
S5/B2	-12.88***	-22.43	-12.23*	-1.06	3.88*	1.7
S5/B3	-12.81***	-47.18**	-12.26*	-1.77*	3.79	2.2
S5/B4	-18.43***	-38.64**	-18.05***	-5.48***	3.11	8.8
S5/B5	-22.17***	-23.87	-28.16***	-4.92***	0.24	8.5
R_M	-17.72***	-19.51	-16.57***	-8.44***	-1.26	6.3
SMB	-3.71	22.28***	-2.21	-0.30	-0.28	4.5
HML	-5.89	19.29	-8.94*	-0.49	-0.27	0.9

2.5.3 Time-Varying Betas in the Three-Factor Model

Conditional asset pricing theory asserts that the significant relation between the predictors and expected returns should disappear when their role as determinant of time-varying risk is recognized. In contrast, the mispricing view argues that their predictive power will persist even when fluctuations in risk are taken into account. Thus, in order to distinguish between both views, we model time variation in risk loadings as a function of the instruments and test whether the alphas produced by conditional specifications of the three-factor model are constant.

Table 2.5 reports regression results for the conditional three-factor model. The table shows the explanatory power of the model in terms of adjusted R^2 . It also provides p -values of F -tests performed to investigate whether the lagged instruments pick up significant variation in risk loadings. The null hypothesis for these tests is that the loadings on the interaction terms between the risk factors and the conditioning variables are jointly equal to zero. p -Values are below 5% for 24 portfolios in the constant alpha case and 22 portfolios when alpha is allowed to vary. The joint Bonferroni test rejects the null hypothesis of constant betas at the 1% level.⁸ Thus, betas exhibit strong time variation, which can be captured by a set of lagged instruments. For most portfolios the adjusted R^2 rises considerably when risk loadings are allowed to fluctuate over time. The explanatory power of the model also increases for many portfolios when time variation in alphas is modeled. This suggests that alphas are not constant over time, even in a model with time-varying betas. We formally test this hypothesis in the next section.

2.5.4 Conditional Alphas in the Three-Factor Model

In Table 2.6 results are shown for tests of the hypothesis that pricing errors are zero and of the weaker hypothesis that alphas are constant through time. We test whether conditional alphas are zero by performing an F -test for the hypothesis that the intercept and the slopes on the lagged instruments are jointly equal to zero. Columns two and three show that the hypothesis of a zero conditional alpha is rejected at the 5% level for 15 portfolios in the constant beta case and 12 portfolios in the conditional three-factor model. The Bonferroni adjusted p -value for a joint test across portfolios is 0.000.

When testing whether alphas are constant, the null hypothesis of the F -test is that the instruments for the conditional alpha can be excluded from the model. Results reported in column four indicate that the weaker hypothesis that alphas are constant is rejected at the 5% level for 15 of the 25 portfolios in a model with constant betas. This implies that the static three-factor model does not adequately explain the dynamics of conditional expected returns. In contrast, column five shows that when betas are allowed to vary the null hypothesis is rejected for only eight portfolios. Thus, the ability of the instruments to predict mispricing diminishes when allowing for time variation in factor loadings. Nevertheless, the joint Bonferroni test still rejects the null hypothesis of constant alphas.

⁸Bonferroni is a multiple-comparison correction for dependence across portfolios.

Table 2.5: **Time-Varying Betas in the Three-Factor Model**

Excess returns on 25 size-B/M portfolios are regressed on a constant, lagged instruments, the three Fama-French factors, and interaction terms between these factors and the instruments. The second column shows the adjusted R^2 for a constant alpha, constant beta model, while the third column presents this statistic for a constant alpha, time-varying beta model. The fourth column reports the p -value of an F -test comparing the R^2 of these two models to test for time-varying betas. $\#p < 0.05$ is the number of p -values below 0.05. *Bonferroni* is the adjusted p -value for a joint test across portfolios of the null hypothesis that betas are constant. The last three columns contain results for a similar analysis but in this case the alphas are allowed to be time-varying.

	Constant Alpha			Time-Varying Alpha		
	Adj. R^2	Adj. R^2	p -Value F -test	Adj. R^2	Adj. R^2	p -Value F -test
	Constant Betas	Dynamic Betas		Constant Betas	Dynamic Betas	
S1/B1	74.0	79.6	0.000	79.3	82.5	0.000
S1/B2	72.6	77.3	0.000	75.3	78.3	0.001
S1/B3	73.3	75.8	0.005	75.5	77.2	0.024
S1/B4	80.5	83.2	0.000	81.9	85.5	0.000
S1/B5	83.9	84.5	0.094	84.6	85.1	0.139
S2/B1	85.1	86.5	0.003	86.5	87.2	0.054
S2/B2	84.4	87.0	0.000	85.0	87.2	0.000
S2/B3	84.6	87.8	0.000	84.8	88.6	0.000
S2/B4	86.1	87.3	0.007	85.8	87.4	0.001
S2/B5	86.0	88.1	0.000	87.1	89.0	0.000
S3/B1	85.5	87.8	0.000	86.3	88.1	0.000
S3/B2	87.7	89.4	0.000	87.5	89.2	0.000
S3/B3	85.5	88.7	0.000	85.7	88.6	0.000
S3/B4	86.1	88.2	0.000	86.4	88.4	0.000
S3/B5	84.2	86.7	0.000	84.6	86.9	0.000
S4/B1	87.7	89.8	0.000	88.4	90.1	0.000
S4/B2	87.8	88.6	0.021	87.9	88.4	0.077
S4/B3	84.7	87.9	0.000	85.3	87.8	0.000
S4/B4	85.6	88.3	0.000	85.8	88.5	0.000
S4/B5	82.5	85.7	0.000	83.2	86.1	0.000
S5/B1	84.9	88.1	0.000	85.0	88.2	0.000
S5/B2	88.9	91.0	0.000	88.8	90.9	0.000
S5/B3	86.2	89.7	0.000	87.0	90.0	0.000
S5/B4	86.1	89.9	0.000	86.6	90.0	0.000
S5/B5	75.7	85.1	0.000	78.9	85.1	0.000
# $p < 0.05$			24			22
Bonferroni			0.000			0.000

In sum, the main conclusion drawn from the time series analysis is that betas are time-varying and that these fluctuations in risk can be picked up by a combination of macroeconomic and portfolio-specific instruments. Conditional specifications of the three-factor model outperform their unconditional counterpart in explaining time variation in expected returns. However, even after taking time variation in betas into account, the model does not fully explain conditional expected returns on the portfolios. Predictable patterns in pricing errors remain, consistent with results documented by Ferson and Harvey (1999) for the United States but contradicting the conclusion of Lewellen (1999) that modeling time variation in risk eliminates the predictive power of B/M.

2.5.5 Cross-Sectional Evidence on the Three-Factor Model

Having found evidence of substantial fluctuations in betas, we now examine whether incorporating time variation in risk is sufficient to eliminate the cross-sectional explanatory power of size, book-to-market, and momentum variables. The cross-sectional analysis is useful to identify the sources of the mispricing detected by the time series regressions. Following Avramov and Chordia (2006a), we evaluate the pricing abilities of alternative model specifications by looking at the significance of Fama-MacBeth parameter estimates for the size, book-to-market, and momentum variables. In addition, we use the time series average of the cross-sectional adjusted R^2 as an informal measure of model performance. In short, a low R^2 and insignificant coefficients can be interpreted as support for the model used to risk-adjust returns, since these imply that the explanatory power of the portfolio characteristics is limited. In contrast, a high adjusted R^2 and significant FM coefficient estimates suggest that size, book-to-market, and momentum effects are not adequately captured by the asset pricing model.⁹

We start off by considering results for FM regressions of raw returns (i.e., *not* adjusted for risk) on a constant and the portfolio characteristics *SIZE*, *B/M*, and the past return variables *RET2-3*, *RET4-6*, and *RET7-12*. Average cross-sectional regression coefficients are reported in column two of Table 2.7 along with their *t*-statistics and the time series average of the monthly adjusted R^2 . The intercept is significant at a 5% level, which shows the presence of time-invariant pricing errors. The coefficient on size is negative and significant. Thus, a size effect is present in the cross-section of European stock returns, which confirms earlier results by Heston, Rouwenhorst, and Wessels (1999). The book-to-market coefficient is positive but insignificant, which means that the value premium is absent. Loadings on all three past return variables are positive and significant. Strong momentum effects in size-B/M portfolios have also been found in the United States by Lewellen (2002). The average adjusted R^2 is 38.0%, which confirms that the characteristics explain a substantial fraction of cross-sectional variation in raw portfolio returns.

⁹A zero R^2 does not necessarily mean that the model completely explains the cross-section of returns. For instance, a significant intercept would imply that the model produces pricing errors unrelated to the characteristics included in the regression.

Table 2.6: **Time-Varying Alphas in the Three-Factor Model**

Excess returns on 25 size-B/M portfolios are regressed on a constant, lagged instruments, the three Fama-French factors, and interaction terms between these factors and the instruments. The second column in this table reports the p -value of an F -test for the hypothesis that the conditional alpha is zero in the constant beta three-factor model while in column three p -values are shown for the same hypothesis when betas are allowed to vary over time as a function of instrumental variables. Columns four and five report p -values for the null hypothesis that alpha is constant in the constant beta three-factor model and a three-factor model with time-varying betas, respectively. $\#p < 0.05$ is the number of p -values below 0.05. *Bonferroni* is the adjusted p -value for a joint test across portfolios.

	Test Zero Alpha		Test Constant Alpha	
	Constant Beta	Dynamic Beta	Constant Beta	Dynamic Beta
S1/B1	0.000	0.000	0.000	0.000
S1/B2	0.000	0.001	0.000	0.019
S1/B3	0.001	0.017	0.001	0.009
S1/B4	0.004	0.000	0.002	0.000
S1/B5	0.024	0.031	0.017	0.038
S2/B1	0.000	0.005	0.000	0.017
S2/B2	0.055	0.296	0.034	0.205
S2/B3	0.124	0.003	0.204	0.004
S2/B4	0.956	0.351	0.923	0.250
S2/B5	0.001	0.001	0.001	0.003
S3/B1	0.007	0.136	0.007	0.084
S3/B2	0.268	0.331	0.832	0.981
S3/B3	0.011	0.012	0.130	0.707
S3/B4	0.154	0.265	0.096	0.179
S3/B5	0.063	0.078	0.061	0.158
S4/B1	0.007	0.127	0.009	0.078
S4/B2	0.044	0.198	0.222	0.768
S4/B3	0.026	0.391	0.020	0.558
S4/B4	0.220	0.304	0.181	0.215
S4/B5	0.020	0.026	0.020	0.062
S5/B1	0.270	0.073	0.313	0.256
S5/B2	0.904	0.804	0.828	0.717
S5/B3	0.009	0.106	0.004	0.064
S5/B4	0.051	0.049	0.039	0.155
S5/B5	0.000	0.031	0.000	0.374
$\#p < 0.05$	15	12	15	8
Bonferroni	0.000	0.000	0.000	0.000

In order to examine whether the high explanatory power and significance of the size and momentum variables persist when cross-sectional differences in risk are taken into account, we risk-adjust returns in first-pass three-factor regressions. FM parameter estimates are shown in column three. The coefficient on size is no longer significant at the 5% level. The intercept has also become insignificant and the adjusted R^2 has fallen sharply to 12.1%. However, all three momentum variables are still significant, confirming the finding of Fama and French (1996) that their three-factor model does not capture the momentum effect.

Results for the conditional Fama-French model are presented in column four. As in the unconditional model, coefficients on the size characteristic and the B/M variable are insignificant. More important, however, is that the loadings on all three past return variables are still significant at the 5% level, suggesting that even a dynamic three-factor model cannot capture the momentum effect. Furthermore, although the conditional model produces the lowest average adjusted R^2 , it still exceeds 10%, reflecting the strong cross-sectional predictive power of the past return variables.

In conclusion, the results presented in Table 2.7 indicate that the Fama-French model eliminates the cross-sectional explanatory power of the size characteristic. Allowing for time variation in risk loadings only leads to a marginal improvement in the pricing ability of the model specifications we consider. In particular, both static and dynamic three-factor models are unable to explain the impact of past returns on the cross-section of portfolio returns.

2.5.6 Business Cycle Risk and Momentum

Chordia and Shivakumar (2002) suggest that in the United States return momentum is related to business cycle variations. In contrast, Griffin, Ji, and Martin (2003) present international evidence that momentum profits cannot be explained by macroeconomic risk factors. However, neither study considers the relation between momentum and macroeconomic state variables in a framework that controls for time-varying risk exposures. Therefore, Avramov and Chordia (2006a) specify a conditional asset pricing model in which alpha is a linear function of business cycle variables. They find that in this model the predictive power of momentum variables becomes insignificant and conclude that a yet undiscovered risk factor related to the business cycle may explain the momentum effect in the United States, supporting the results of Chordia and Shivakumar (2002).

It is interesting to see whether we can confirm these findings using the 25 European size-B/M portfolios as test assets. Following the approach of Avramov and Chordia (2006a), we allow alpha to vary over time as a function of three business cycle variables: risk-free rate, default spread, and term spread. The dependent variable in the cross-sectional regressions is then α_{i0} , the part of mispricing unrelated to the business cycle.

The FM estimates in columns five and six of Table 2.7 show the relation between the portfolio characteristics and the part of risk-adjusted returns unrelated to business cycle effects. The results in these columns indicate that all momen-

Table 2.7: Fama-MacBeth Regressions of Raw Returns and Three-Factor Risk-Adjusted Returns

This table reports Fama-MacBeth coefficient estimates. In the second column the dependent variable is the excess portfolio return unadjusted for risk. In the third column the dependent variable is the excess return risk-adjusted using the Fama-French three-factor model with constant alphas and betas. In the fourth column the dependent variable is the excess portfolio return risk-adjusted using a conditional three-factor model with constant alphas, where factor loadings are scaled by the default spread, size, book-to-market, and interaction terms between default spread and size and between default spread and book-to-market,

$$\beta_{ikt} = \beta_{ik1} + \beta_{ik2}DEF_t + (\beta_{ik3} + \beta_{ik4}DEF_t)SIZE_{it} + (\beta_{ik5} + \beta_{ik6}DEF_t)B/M_{it},$$

where DEF is the default spread and $SIZE$ and B/M are the logarithm of portfolio market capitalization in billions of euros and the logarithm of portfolio book-to-market, respectively. In the fifth and sixth column alphas are time-varying and a function of macroeconomic variables,

$$\alpha_{it} = \alpha_{i0} + \alpha_{i1}DEF_t + \alpha_{i2}R_{Ft} + \alpha_{i3}TERM_t,$$

where R_F is the risk-free rate and $TERM$ is the term spread. The dependent variable in the fifth and sixth column is α_{i0} , the part of the risk-adjusted return that is unrelated to business cycle effects.

Portfolio characteristics are the explanatory variables in the cross-sectional regressions. $SIZE$ and B/M are expressed as logarithms of market capitalization and book-to-market, respectively, and $RET2-3$, $RET4-6$, and $RET7-12$ are cumulative past returns. All five characteristics are measured as deviations from their cross-sectional mean in each month. The sample period is February 1986 through June 2002. Adj. R^2 is the time series average of the monthly adjusted R^2 . t -Statistics are reported in parentheses.

	Return	Constant Alpha		Dynamic Alpha	
		Constant β	Dynamic β	Constant β	Dynamic β
Intercept	0.776 (2.190)	0.077 (0.879)	0.057 (0.705)	0.211 (1.864)	0.132 (1.621)
SIZE	-0.093 (-1.970)	-0.027 (-0.978)	-0.024 (-1.027)	-0.019 (-0.536)	-0.014 (-0.520)
B/M	0.006 (0.051)	-0.066 (-1.106)	-0.005 (-0.084)	0.013 (0.307)	0.040 (0.702)
RET2-3	0.044 (2.934)	0.036 (2.372)	0.040 (2.909)	0.042 (3.019)	0.045 (3.457)
RET4-6	0.029 (2.308)	0.028 (2.164)	0.027 (2.293)	0.029 (2.278)	0.029 (2.434)
RET7-12	0.044 (5.623)	0.041 (5.111)	0.036 (4.920)	0.040 (5.034)	0.036 (4.860)
Adj. R^2	38.0	12.1	10.6	12.0	10.5

tum variables remain significant, irrespective of whether betas are constant or time-varying. Moreover, the cross-sectional predictive power of the characteristics, measured by the adjusted R^2 , is almost identical to the values reported in columns three and four. In sum, we find no evidence that momentum is related to business cycle variations in a way that can be captured by the conditional three-factor model we consider. These findings support the conclusion of Griffin, Ji, and Martin (2003): outside the United States, macroeconomic factors do not explain return momentum.

2.5.7 Robustness of Empirical Results

As a robustness check, we purge the 25 size-B/M portfolios from possible country and sector effects to determine whether our results are affected by country or sector tilts in the portfolios. We control for country and sector effects by first performing cross-sectional regressions of firm size and B/M on country and sector dummies,

$$x_{it} = \kappa_t + \sum_{j=1}^{J-1} \theta_{jt} S_{ij} + \sum_{h=1}^{H-1} \psi_{ht} C_{ih} + \tau_{it}, \quad (2.8)$$

where x_{it} is a vector that contains the size and book-to-market characteristics of firm i at date t , S_{ij} a dummy variable equal to one if firm i belongs to sector j and zero otherwise, and C_{ih} a dummy variable that equals one if firm i is listed in country h . By leaving out the sector (country) dummies we can purge the characteristics from country (sector) effects only. Subsequently, the vector of residuals τ_{it} from (2.8) is used to sort the stocks into 25 size-B/M portfolios.

The bottom half of Table 2.2 reports value-weighted returns for the 25 size-B/M portfolios constructed using the country- or sector-neutral characteristics. Similar to the original portfolios shown in the upper left corner, the purged portfolios exhibit a significant size effect but insignificant value premium, although within some individual size quintiles a significant B/M effect can be observed. In general, however, returns on the country-neutral and sector-neutral portfolios do not deviate strongly from the returns on the original portfolios. We can confirm that using these portfolios as test assets in the empirical analysis does not alter our conclusions.¹⁰

The second check on our results deals with possible correlation between errors in the factor loadings estimated in the first-pass time series regressions and the non-risk portfolio characteristics used as explanatory variables in the second-stage cross-sectional regressions. Although the factor loadings are correlated with the portfolio characteristics included in P_t in equation (2.7), Brennan, Chordia, and Subrahmanyam (1998) argue that there is no a priori reason to believe that the *errors* in the estimated loadings are correlated with the characteristics. Nevertheless, if they are not independent, the cross-sectional regression coefficients on the

¹⁰Results for all robustness checks are omitted in the interest of parsimony and available upon request.

portfolio characteristics will be correlated with the factor returns and consequently the standard Fama-MacBeth estimator will be biased.

Therefore, Brennan, Chordia, and Subrahmanyam (1998) propose to calculate a purged estimator for each of the characteristics. This estimator is unbiased when estimation errors in the factor loadings are correlated with the characteristics in P_t , provided that the factor premia are serially uncorrelated. The purged estimator is the intercept in an OLS regression of the original monthly cross-sectional parameter estimates on a constant and the time series of factor realizations FF_{kt} . It turns out that our empirical results are almost unchanged when the purged estimator is used instead of the standard FM estimator.

A third check is motivated by Shanken (1992), who points out that the Fama-MacBeth procedure overstates the precision of parameter estimates in the second-stage cross-sectional regressions by ignoring estimation errors in factor loadings obtained from first-pass time series regressions. Shanken suggests a solution for this problem that explicitly adjusts the standard errors for measurement errors, assuming conditional homoskedasticity of returns. Applying this correction leads to t -statistics that are only slightly lower than the standard OLS t -statistics shown in Table 2.7. Hence, our conclusions from the cross-sectional analysis still hold.

Fourth, instead of estimating the time series regressions by ordinary least squares (OLS) regressions, we also perform seemingly unrelated regressions (SUR) to account for contemporaneous correlation in residuals across portfolios. However, estimation results are very similar and all main conclusions remain unchanged.

Finally, we repeat the empirical analysis for the CAPM and the Carhart (1997) four-factor model, which adds a momentum factor to the three-factor model. As expected, the three-factor model is clearly superior to the CAPM in terms of time series and cross-sectional explanatory power. Although we find evidence of significant time variation in CAPM betas, a conditional CAPM is not able to capture size and momentum effects in portfolio returns. These results are consistent with findings documented by Lewellen and Nagel (2006) for the United States, who conclude that allowing for time variation in beta does little to salvage the CAPM. Results for the Carhart four-factor model are very similar to those reported for the Fama-French model. Most importantly, the Carhart model also fails to explain the strong momentum effects in the 25 size-B/M portfolios.

2.6 Conclusion

This chapter examines whether risk loadings in the Fama and French (1993) three-factor model are time-varying and if so, to what extent conditional specifications of the model can eliminate well-known anomalies in European stock markets. Our work is motivated by mixed empirical evidence on the performance of conditional asset pricing models in the United States. Prior research shows that several firm characteristics like size, book-to-market, and past returns have explanatory power for the cross-section of returns. Furthermore, it has been found that size, B/M, and macroeconomic variables predict time variation in expected returns.

We combine these findings to examine whether the predictive power of these variables is due to time-varying risk, as suggested by conditional asset pricing theory, or to mispricing. Specifically, we test the ability of static and dynamic specifications of the three-factor model to price 25 size-B/M portfolios. The portfolios are constructed using merged data from 16 European stock markets, thereby providing out-of-sample evidence on conditional asset pricing models.

We identify important differences between U.S. and European data. In particular, while the size effect has vanished in the United States after its discovery, we show that it is still present in Europe. Moreover, in contrast to U.S. evidence, we find that the notoriously hard to price small-growth portfolio displays significant *positive* pricing errors in European markets. We also show that a set of macroeconomic and portfolio-specific variables has substantial predictive power for the European size-B/M portfolios. Our time series tests reveal that these variables pick up time variation in risk. As a result, conditional specifications of the three-factor model outperform their static counterpart in explaining time variation in expected returns. Nevertheless, even after allowing for fluctuations in factor loadings, pricing errors for some portfolios are still significant and predictable.

In order to identify the sources of mispricing, we apply the cross-sectional testing framework of Avramov and Chordia (2006a). While the three-factor model captures the size effect, both static and dynamic specifications of the model fail to eliminate the strong cross-sectional predictive power of momentum variables. Conditioning does little to improve the cross-sectional pricing ability of the model. Finally, using European data we find no empirical support for the hypothesis put forward by Avramov and Chordia (2006a) that momentum is related to business cycle variations. Thus, although the evidence of time-varying risk motivates the use of conditional asset pricing models, more is needed to revive modern finance.

Chapter 3

Efficient Estimation of Firm-Specific Betas and its Benefits for Asset Pricing Tests and Portfolio Choice*

This chapter improves both the specification and estimation of firm-specific betas. Time variation in betas is modeled by combining a parametric specification based on economic theory with a nonparametric approach based on data-driven filters. We increase the precision of individual beta estimates by setting up a hierarchical Bayesian panel data model that imposes a common structure on parameters. We show that these accurate beta estimates lead to a large increase in the cross-sectional explanatory power of the conditional CAPM. Using the betas to forecast the covariance matrix of stock returns also results in a significant improvement in the out-of-sample performance of minimum variance portfolios.

3.1 Introduction

Precise estimates of firm-specific betas are crucial in many applications of modern finance theory, including asset pricing, corporate cost-of-capital calculations, and risk management. For instance, portfolio managers often have to ensure that their risk exposure stays within predetermined limits and managers need estimates of their company's beta to make capital budgeting decisions. Academics and practitioners have taken two approaches to estimating betas. The first method groups stocks into portfolios to reduce measurement error, assuming that all stocks within a given portfolio share the same beta (e.g., Fama and MacBeth (1973)). The second approach consists of estimating separate time series regressions for

*This chapter is based on Cosemans, Frehen, Schotman, and Bauer (2009).

all firms to obtain individual betas (e.g., Brennan, Chordia, and Subrahmanyam (1998)).

Apart from this lack of consensus in the literature about the best method to estimate betas, existing studies also fail to provide clear guidance on the best way to model betas. Although a large body of empirical evidence suggests that betas vary over time, existing work uses different specifications to model these changes in betas.¹ Many studies use a parametric approach proposed by Shanken (1990), in which variation in betas is modeled as a linear function of conditioning variables. An alternative, nonparametric approach to model risk dynamics is based on purely data-driven filters, including short-window regressions (Lewellen and Nagel (2006)) and rolling regressions (Fama and French (1997)).²

In this chapter we improve both the *specification* and *estimation* of time-varying, firm-specific betas. We improve the specification of betas by combining the parametric and nonparametric approaches to modeling time variation in betas. Because the key strengths of each approach are the most important weaknesses of the other, we argue and show that a combination of the two methods leads to more accurate betas than those obtained from each of the two approaches separately. We allow the optimal mix of the two methods to vary across stocks, since individual firms may benefit more or less from either specification, and over time, because the preferred combination during stable market conditions may be different from that in turbulent time periods.

The parametric specification is appealing from a theoretical perspective because it explicitly links time variation in betas to macroeconomic state variables and firm characteristics (see, e.g., Gomes, Kogan, and Zhang (2003)). However, the main drawback of this approach is that the investor's set of conditioning information is unobservable. Ghysels (1998) shows that misspecifying beta risk may result in pricing errors that are even larger than those produced by an unconditional asset pricing model. In addition, this method can produce excessive variation in betas due to sudden spikes in the macroeconomic variables that are often used as instruments. Finally, many parameters need to be estimated when a large number of conditioning variables is included, which leads to noisy estimates for stocks with only a limited number of time series observations. An important advantage of the nonparametric approaches is that they preclude the need to specify conditioning variables, which makes them more robust to misspecification. However, the time series of betas produced by a data-driven approach will always lag the true variation in beta, because using a window of past returns to estimate the beta at a given point in time gives an estimate of the average beta during this time period. Although reducing the length of the window results in timelier betas, the estimation precision of these betas will also decrease.

¹See, for instance, Jagannathan and Wang (1996), Lewellen (1999), Ferson and Harvey (1999), Lettau and Ludvigson (2001b), Andersen, Bollerslev, Diebold, and Wu (2005), Avramov and Chordia (2006a), Ang and Chen (2007), and Ang and Kristensen (2009).

²An alternative approach has been proposed by Chang, Christoffersen, Jacobs, and Vainberg (2009), who calculate forward-looking betas using the information embedded in option data. A drawback of this method is that it requires a cross-section of liquid stock options, which is not available for many small firms.

We improve the estimation of individual stock betas by setting up a Bayesian panel data model that exploits the information in the cross-section of firms to obtain more precise estimates. In particular, we specify hierarchical prior distributions that impose a common structure on parameters while still allowing for cross-sectional heterogeneity. Bayesian methods are especially attractive in settings with individual-level heterogeneity in multiple parameters, because only the parameters of the hierarchical priors where the parameters are assumed to be drawn from have to be estimated. In contrast, methods that estimate every parameter individually without linking it to similar parameters, such as estimating a separate time series regression for every single firm, suffer from poor estimation precision, particularly when the number of time series observations is limited. Intuitively, the Bayes estimator can be interpreted as a weighted average of the least squares estimator for a given firm and the cross-sectional average coefficient. In particular, it shrinks the least squares estimator of the firm-specific parameters towards the cross-sectional mean. When the number of data points increases, the weight gradually shifts from the prior to the data.

Our panel data approach uses both daily returns and monthly firm characteristics to capture the cross-sectional heterogeneity and time series dynamics of monthly betas. Including cross-sectional information increases the accuracy of firm-specific betas because previous studies document a strong cross-sectional relationship between beta and firm characteristics (see, e.g., Fama and French (1992)). Existing work further shows that the use of high-frequency returns yields more precise and timelier estimates of beta than using monthly returns (see, e.g., Bollerslev and Zhang (2003)). We use daily returns instead of intraday returns because market microstructure frictions put an upper limit on the frequency that can be used to estimate betas in practice. We combine the data sampled at different frequencies by implementing the mixed data sampling (MIDAS) approach of Ghysels, Santa-Clara, and Valkanov (2005), which determines the optimal weights given to past data.

We estimate the model using a large panel of individual stocks, which offers several advantages over the alternative of sorting stocks into portfolios based on characteristics. First, aggregating stocks into portfolios may conceal important information contained in individual stock betas. Ang, Liu, and Schwarz (2008) show that risk premia can be estimated more precisely using individual stocks instead of portfolios, because creating portfolios reduces the cross-sectional variation in betas. A second important drawback is that due to the strong factor structure in the 25 size-B/M sorted portfolios that are often used as test assets in asset pricing studies, standard cross-sectional tests are flawed and have low power to reject a model, as shown by Lewellen, Nagel, and Shanken (2009). Third, when stocks are grouped into portfolios based on characteristics that have been identified by previous research as determinants of average returns instead of being based on economic theory, the evidence against asset pricing models may be overstated because of data-snooping biases (Lo and MacKinlay (1990)).

Despite the benefits of using individual stocks, most asset pricing studies use characteristics sorted portfolios because it is difficult to estimate firm-level pa-

rameters with a reasonable degree of precision. Notable exceptions are Brennan, Chordia, and Subrahmanyam (1998) and Avramov and Chordia (2006a), who use a two-stage approach to study the impact of characteristics on risk-adjusted returns. However, both studies estimate a separate time series regression for each firm, which leads to imprecise beta estimates, particularly for firms with a short return history. Fama and French (2008) even conclude that “given the imprecision of beta estimates for individual stocks, little is lost in omitting them from the cross-section regressions.”

Our main empirical findings are as follows. First, we show that modeling time-varying betas as a function of both conditioning variables and past returns dominates traditional specifications in which betas depend on only one of these components. Combining the two specifications produces superior beta estimates because they capture different aspects of beta dynamics. We also find that the optimal mix of these specifications varies both over time and across stocks. Second, we show that our panel data approach produces more accurate estimates of firm-specific betas than those obtained from the traditional approach of estimating time series regressions. Specifically, for the average firm the posterior standard deviation of beta is significantly larger in a time series regression than in the panel model. Third, we document strong cross-sectional heterogeneity in firm betas within the 25 size-B/M portfolios that are commonly used to test asset pricing models. This confirms that aggregating stocks into portfolios conceals important information contained in individual stocks and shrinks the cross-sectional variation in betas.

We demonstrate that a better specification and more precise estimation of firm betas has important benefits for asset pricing tests and portfolio choice. In particular, we show that the betas generated by our model have significant explanatory power for the cross-section of returns. Using stocks as test assets and estimating betas in a panel model results in more efficient parameter estimates in cross-sectional asset pricing tests than using portfolios. The estimate of the market premium is positive and statistically significant, even after controlling for firm characteristics. We illustrate the value of our beta specification and estimation method for portfolio choice by using the betas to forecast the covariance matrix of stock returns. We find that the global minimum variance portfolio that is formed using this covariance matrix outperforms minimum variance portfolios based on other strategies, including the naive $1/N$ rule, the sample covariance matrix, and a static one-factor model for estimating covariances.

The chapter proceeds as follows. In Section 3.2 we introduce our specification of time-varying individual stock betas in a panel data framework. Section 3.3 explains the Bayesian approach to inference and Section 3.4 describes the data. We report our empirical results in Section 3.5 and discuss the asset pricing and portfolio choice applications in Section 3.6. Section 3.7 concludes.

3.2 The Model

In this section we describe our model for individual betas. For simplicity, we discuss our approach in a conditional CAPM setting but it is straightforward to extend our work to multifactor models. Our goal is to show how to improve the specification and estimation of time-varying, firm-specific betas in any factor model. We start from the following panel data model for excess returns on individual stocks,

$$r_{it} = \alpha_i + \beta_{it-1}r_{Mt} + \epsilon_{it}, \quad (3.1)$$

where r_{it} is the excess return on stock i in month t , α_i is the risk-adjusted return, β_{it-1} is the conditional market beta, r_{Mt} is the excess market return, and ϵ_{it} is a zero-mean, normally distributed idiosyncratic return shock. Following Avramov and Chordia (2006b), we assume that the covariance matrix of these shocks is diagonal and that idiosyncratic volatility is constant.

Our specification of the conditional beta consists of two components: one part is the realized beta, b_{it} , and the other part is the fundamental beta, β_{it}^* ,

$$\beta_{it} = \phi_{it}b_{it} + (1 - \phi_{it})\beta_{it}^*, \quad (3.2)$$

where ϕ_{it} and $(1 - \phi_{it})$ measure the proportion of the beta of firm i that is explained by the realized beta and fundamental beta, respectively. Hereafter we refer to this mixture of realized and fundamental betas as the *mixed* beta. We allow the optimal combination of fundamental and realized betas to vary not only across firms but also over time. Time variation in ϕ_{it} is modeled as a linear function of market variance, because the best mix of fundamental and realized betas in turbulent market conditions can be very different from that in stable periods,

$$\phi_{it} = \phi_{0i} + \phi_{1i}V_{Mt}, \quad (3.3)$$

where V_{Mt} is the realized market variance, which we calculate by summing the squared daily market returns over the past year.³ We take the logarithm of market variance to reduce the impact of outliers and then subtract its time series average and divide by its standard deviation, so that it has mean zero and standard deviation equal to one.

b_{it} is the realized beta that we estimate using daily data according to the Mixed Data Sampling (MIDAS) approach introduced by Ghysels, Santa-Clara, and Valkanov (2005). We choose to estimate realized betas using daily returns because these provide a reasonable balance between efficiency and robustness to microstructure noise (see, Campbell, Lo, and MacKinlay (1997)). However, even at a daily frequency the betas of less liquid stocks might be biased downward.

³An interesting avenue for future research would be to make the weight ϕ_{it} a function of the relative precision of the fundamental and realized beta estimates, as measured by their posterior variance. Our model could then be interpreted as a generalization of the shrinkage approach to estimating security betas developed by Vasicek (1973), as it allows for time variation in beta and shrinks the data-driven realized beta towards the theoretically motivated fundamental beta.

Following Scholes and Williams (1977), we therefore control for nonsynchronous trading effects by adding the covariance of the stock's return with the one-day lagged market return to the model.

The MIDAS approach differs from traditional rolling window estimators of betas by selecting the optimal window for estimating betas using a flexible weighting function. Ghysels, Santa-Clara, and Valkanov (2005) use the MIDAS approach to estimate the market's conditional variance and find that it is superior to traditional GARCH and rolling window methods. In particular, our MIDAS estimator of realized betas is given by,

$$b_{it} = \frac{\sum_{\tau=1}^{\tau^{max}} w_{t-\tau} r_{it-\tau}^{(d)} r_{Mt-\tau}^{(d)}}{\sum_{\tau=1}^{\tau^{max}} w_{t-\tau} r_{Mt-\tau}^{(d)} r_{Mt-\tau}^{(d)}} + \frac{\sum_{\tau=1}^{\tau^{max}} w_{t-\tau} r_{it-\tau}^{(d)} r_{Mt-\tau-1}^{(d)}}{\sum_{\tau=1}^{\tau^{max}} w_{t-\tau} r_{Mt-\tau-1}^{(d)} r_{Mt-\tau-1}^{(d)}}, \quad (3.4)$$

where t refers to a particular month, τ to a particular trading day, and $w_{t-\tau}$ to the weight given to the product of the return on stock i and the market return, $r_{it-\tau}^{(d)} r_{Mt-\tau}^{(d)}$, and to the squared market return, $r_{Mt-\tau}^{(d)} r_{Mt-\tau}^{(d)}$, on day $t - \tau$. We set the maximum window length τ^{max} equal to 250 trading days, which corresponds to one year of trading.

We parameterize the weights as a beta function,

$$w_{t-\tau} = \frac{f\left(\frac{\tau}{\tau^{max}}, \kappa_1; \kappa_2\right)}{\sum_{\tau=1}^{\tau^{max}} f\left(\frac{\tau}{\tau^{max}}, \kappa_1; \kappa_2\right)}, \quad (3.5)$$

where $f(\frac{\tau}{\tau^{max}}, \kappa_1; \kappa_2)$ is the density of a beta distribution. As pointed out by Ghysels, Santa-Clara, and Valkanov (2005), the specification based on the beta function has several advantages. First, it ensures that the weights are positive and sum to one. Second, it is parsimonious because only two parameters need to be estimated. Third, it is flexible as it can take various shapes for different values of the two parameters. We impose a downward sloping pattern on the weights by setting κ_1 equal to 1, which further reduces the number of parameters that need to be estimated. $\kappa_1 = \kappa_2 = 1$ implies equal weights, which corresponds to a rolling window estimator of beta on daily data. $\kappa_1 = 1$ and $\kappa_2 > 1$ correspond to the case of decaying weights. In general, the higher κ_2 , the faster the rate of decay and the quicker beta responds to new information.

β_{it}^* is the fundamental beta, parameterized as a linear function of conditioning variables,

$$\beta_{it}^* = \delta_0 + \delta_1' [Z_{it} \otimes BC_t], \quad (3.6)$$

where Z_{it} is a vector that contains L firm characteristics and BC_t is a vector that contains a constant and M business cycle variables. This specification allows the relation between beta and firm characteristics to vary over the business cycle. Modeling beta dynamics as a linear function of a set of predetermined instruments goes back to Shanken (1990) and is consistent with the economic motivation for conditional asset pricing models, in which the stochastic discount factor is a function of macroeconomic state variables and factor premia.

We include both firm-specific and macroeconomic variables as instruments for fundamental betas because of their documented predictive power for returns (Fama and French (1989) and Lewellen (1999)). Empirical evidence that systematic risk is related to firm characteristics and business cycle variables is provided by, among others, Jagannathan and Wang (1996), Lettau and Ludvigson (2001b), Avramov and Chordia (2006a), and Goetzmann, Watanabe, and Watanabe (2009). The theoretical motivation for choosing firm characteristics as instruments is given by Berk, Green, and Naik (1999) and Gomes, Kogan, and Zhang (2003), who show that the ability of size and book-to-market to explain the cross-section of returns is due to their correlation with the true conditional market beta. They decompose firm value into the value of assets in place and the value of growth options and demonstrate that size captures the component of a firm's systematic risk related to its growth options whereas the book-to-market ratio is a measure of the risk of the firm's assets in place. Zhang (2005) extends this work and argues that because of costly reversibility of capital, value firms have countercyclical betas while betas of growth stocks are procyclical. Because the price of risk is also countercyclical his model can explain the value premium within a rational framework. In addition to size and B/M, we also select firm-specific momentum as a conditioning variable to examine whether the momentum effect is related to beta dynamics. Theoretical support for including macroeconomic variables is provided by Santos and Veronesi (2004), who show within a general equilibrium model that market betas vary substantially with the business cycle. Our choice of business cycle variables is motivated by previous work (e.g., Ferson and Harvey (1999)) and includes the default spread, dividend yield, one-month T-bill rate, and term spread.

Substituting (3.2) and (3.6) into (3.1) leads to the following specification

$$r_{it} = \alpha_i + \phi_{it}b_{it-1}r_{Mt} + (1 - \phi_{it})(\delta_0 + \delta'_1[Z_{it-1} \otimes BC_{t-1}])r_{Mt} + \epsilon_{it}. \quad (3.7)$$

A key objective in this chapter is to determine whether the time series dynamics and cross-sectional variation in betas is better described by lagged firm characteristics and macroeconomic state variables, by past realized betas, or by a linear combination of both. Therefore, we are primarily interested in the parameter ϕ_{it} and compare three different specifications: (1) mixed beta (ϕ_{it} unrestricted), (2) fundamental beta ($\phi_{it} = 0$), and (3) realized beta ($\phi_{it} = 1$).

3.3 Methodology

3.3.1 Bayesian Methods

We estimate the model parameters using Bayesian methods.⁴ The main advantage of Bayesian inference in our setting is that it allows a very flexible specification for

⁴Bayesian methods have been used in a number of asset pricing studies, including Shanken (1987), Harvey and Zhou (1990), Kandel, McCulloch, and Stambaugh (1995), and Cremers (2006). These papers focus on portfolios and assume that betas are constant. Ang and Chen (2007) and Jostova and Philipov (2005) use Bayesian techniques to obtain time-varying portfolio betas, which they model as latent autoregressive processes.

describing the dynamics in beta by imposing a common structure on the model parameters. Updating beliefs according to Bayes' theorem implies that the joint posterior density of the parameters, $p(\theta|y)$, is proportional to the likelihood times the prior density,

$$p(\theta|y) \propto p(y|\theta)p(\theta), \quad (3.8)$$

where θ is the set of all parameters and y is the full set of data.

The likelihood function for the model in equation (3.7) is given by

$$p(y|\theta) = \prod_{i=1}^N \prod_{t \in T_i} (\sigma_{\epsilon_i}^2)^{-\frac{1}{2}} \exp \left[-\frac{1}{2\sigma_{\epsilon_i}^2} (r_{it} - \alpha_i - \beta_{it-1} r_{Mt})^2 \right], \quad (3.9)$$

where $\sigma_{\epsilon_i}^2$ is the idiosyncratic return variance, β_{it-1} is defined as in equation (3.2), N is the number of stocks in the sample, and T_i is the number of monthly return observations for firm i .

3.3.2 Prior Distributions

We specify conditionally conjugate, independent, hierarchical priors that impose a common structure on the model parameters while still allowing parameters to vary across firms. Thus, our setup combines the benefits of a portfolio approach to estimating betas (e.g., Fama and MacBeth (1973)) and an approach in which a separate regression is estimated for each firm (e.g., Avramov and Chordia (2006a)). Specifically, our choice of prior distributions is as follows

$$\begin{aligned} \alpha_i &\sim N(0, \sigma_\alpha^2) && \text{with } \sigma_\alpha^2 \sim IG(0.001, 0.001), \\ \phi_{0i} &\sim N(0.5, \sigma_{\phi_0}^2) && \text{with } \sigma_{\phi_0}^2 \sim IG(0.001, 0.001), \\ \phi_1 &\sim N(0, \sigma_{\phi_1}^2) && \text{with } \sigma_{\phi_1}^2 \sim IG(0.001, 0.001), \\ \delta_0 &\sim N(0, \sigma_{\delta_0}^2) && \text{with } \sigma_{\delta_0}^2 \sim IG(0.001, 0.001), \\ \delta_1 &\sim N(0, \Omega_{\delta_1}) && \text{with } \Omega_{\delta_1}^{-1} \sim Wish([(L + LM)I]^{-1}, (L + LM)), \\ \sigma_{\epsilon_i}^2 &\sim IG(0.001, 0.001). \end{aligned}$$

We use diffuse priors to minimize their influence on the posterior densities. Following Jostova and Philipov (2005), we specify noninformative prior distributions for the variance parameters σ_α^2 , $\sigma_{\phi_0}^2$, $\sigma_{\phi_1}^2$, $\sigma_{\delta_0}^2$, and the idiosyncratic variance $\sigma_{\epsilon_i}^2$, by setting the scale and shape parameters A and B of their inverse gamma (IG) prior distributions equal to 0.001. We set the degrees of freedom parameter ψ of the Wishart prior for $\Omega_{\delta_1}^{-1}$ equal to the dimension of this matrix, $(L + LM)$, because this value gives the lowest possible weight to prior information (see Gelman, Carlin, Stern, and Rubin (2004)). We set the scale matrix of the Wishart prior equal to $[(L + LM)I]^{-1}$, so that the prior mean of $\Omega_{\delta_1}^{-1}$ is equal to the identity matrix. We give equal prior weight to the fundamental beta and the realized beta by setting the prior mean of ϕ_{0i} equal to 0.5 and the prior mean of ϕ_1 equal to

^{0.5} We parameterize the MIDAS weights as a beta function and set κ_1 equal to 1. To rule out cases where more recent returns receive less weight than observations in the more distant past, i.e., when $\kappa_2 < 1$, we constrain κ_2 to the interval $[1, 26]$. When $\kappa_2 = 1$ all 250 daily returns receive equal weight in the estimation and when $\kappa_2 = 26$ the cumulative weight given to the 40 most recent days is 99%. We implement this restriction by a change of variable, $\kappa_2 = 1 + 25\kappa_2^*$. For κ_2^* we choose a uniform prior, $\kappa_2^* \sim U[0, 1]$.

3.3.3 Bayesian Inference

We employ Markov Chain Monte Carlo (MCMC) methods to sample from the joint posterior distribution of all parameters in θ . The basic idea is to construct a Markov chain such that the chain converges to a unique stationary distribution that is the posterior density, $p(\theta|y)$. We use the Gibbs sampler, which involves the sequential drawing from the full conditional posterior densities, to obtain draws from the joint posterior density. In particular, first the set of parameters θ is partitioned into B blocks $(\theta_1, \theta_2, \dots, \theta_B)$. At each iteration of the Gibbs sampler each block is sampled from its posterior distribution conditional on all other blocks and the data. Because the conditional posterior density of κ_2 has a nonstandard form, we cannot directly sample from it. Therefore, we use the Metropolis-Hastings algorithm, in which candidate parameter values are drawn from a proposal density and accepted with a certain probability that is highest in areas of the parameter space where the posterior density is highest (see Chib and Greenberg (1995)). Details on the derivation of the joint posterior density and the conditional posterior distributions are provided in Appendix 3.A.

Iterations of the chain converge to draws from the joint posterior. We check convergence by inspecting the standardized cumsum statistics, as suggested by Bauwens, Lubrano, and Richard (1999), applying the partial means test based on numerical standard errors, explained by Geweke (2005), and calculating the Gelman-Rubin statistic that compares the variation in output between and within chains, described by Gelman, Carlin, Stern, and Rubin (2004). These diagnostics indicate that the parameter chains have converged after 1,000 iterations. In our empirical analysis we therefore run 5,000 iterations and discard the first 1,000 iterations as burn-in period. The remaining draws are used to summarize the posterior density and to conduct inference.

3.4 Data

The firm data comes from CRSP and Compustat and consists of the monthly return, size, and book-to-market value of all NYSE- and AMEX-listed stocks. To calculate realized betas we also retrieve daily returns of these stocks from CRSP. The sample covers the period from July 1964 to December 2006. Following

⁵We also considered specifications with μ_{ϕ_0} set equal to 0 or 1. Results for these priors (available upon request) show that our findings are robust to the choice of this prior mean.

Avramov and Chordia (2006a), we include a stock in the analysis for a given month t if it satisfies the following criteria. First, its return in the current month t and in the previous 36 months has to be available. Second, data should be available in month $t-1$ for size as measured by market capitalization and for the book-to-market ratio. We calculate the book-to-market ratio using accounting data from Compustat as of December of the previous year. Finally, in line with Fama and French (1993), we exclude firms with negative book-to-market equity. Imposing these restrictions leaves a total of 5,017 stocks over the full sample period and an average of 1,815 stocks per month.

Table 3.1 presents summary statistics for the data set. Panel A reports the mean, median, standard deviation and 5th, 25th, 75th, and 95th percentile values of the distribution of excess stock returns and firm characteristics across all data points. The average monthly excess stock return is 0.69% and the median is -0.16%. The mean (median) firm size is \$1.59 (\$0.16) billion. Because the distribution of the book-to-market ratio contains some extreme values, we trim all book-to-market outliers to the 0.5th and 99.5th percentile values of the distribution. After trimming, the average (median) book-to-market ratio equals 0.96 (0.75). The cumulative return over the twelve months prior to the current month, which we use as a proxy for momentum, has a mean of 14.65% and a median of 8.60%. Because the distributions of firm size and book-to-market display considerable skewness, we use the logarithmic transformations of these variables in the analysis. Furthermore, we normalize the characteristics by expressing them as deviations from their cross-sectional means to remove any time trend in the average value of the characteristics.

We further retrieve data for the four macroeconomic variables that we use as instruments for the fundamental beta, i.e., the default spread, dividend yield, one-month Treasury bill rate, and term spread. We define the default spread as the yield differential between bonds rated Baa by Moody's and bonds with a Moody's rating of Aaa. The dividend yield is calculated as the sum of the dividends paid on the value-weighted CRSP index over the previous twelve months divided by the current level of the index. The term spread is defined as the yield difference between ten-year and one-year Treasury bonds. Panel B shows descriptive statistics for the macroeconomic variables. The average default spread is 1.02%, the mean dividend yield equals 3.01%, the average one-month T-bill rate is 5.69%, and the average term spread is 0.85%.

3.5 Empirical Results

In Section 3.5.1 we study whether betas are driven by lagged conditioning variables or by past realized betas. Section 3.5.2 compares the efficiency of the beta estimates produced by the panel model to that of those obtained from time series regressions. In Section 3.5.3 we illustrate the loss of information from aggregating stocks into portfolios by showing the cross-sectional variation in firm-level betas within the 25 size-B/M portfolios that are often used to test asset pricing models.

Table 3.1: **Summary Statistics for Firm Characteristics and Macroeconomic Variables**

This table presents descriptive statistics for stock returns, firm characteristics, and macroeconomic variables for 510 months from July 1964 through December 2006. Panel A reports the mean, median, standard deviation and 5th, 25th, 75th, and 95th percentile values of firm characteristics for a total of 5,017 stocks over the full sample period and an average of 1,815 stocks per month. We include a stock in the sample for a given month t if it satisfies the following criteria. First, its return in the current month t and in the past 36 months has to be available. Second, data should be available in month $t-1$ for size as measured by market capitalization and for the book-to-market ratio. We exclude firms with negative book-to-market equity. $XRET$ is the return in excess of the risk-free rate and MV represents the market capitalization in billions of dollars. BM is the book-to-market ratio, for which values smaller than the 0.5th percentile and values greater than the 99.5th percentile of the distribution are set equal to the 0.5th percentile and 99.5th percentile values, respectively. MOM is the cumulative return over the twelve months prior to the current month. Panel B shows the mean, median, standard deviation and 5th, 25th, 75th, and 95th percentile values of the distribution of a set of macroeconomic variables. DEF is the default spread, defined as the yield differential between bonds rated Baa by Moody's and bonds with a Moody's rating of Aaa. DY is the dividend yield on the value-weighted CRSP index, calculated as the sum of the dividends paid on the index over the previous twelve months divided by the current level of the index. $TBILL$ is the one-month Treasury bill rate. $TERM$ is the term spread, defined as the yield difference between ten-year and one-year Treasury bonds.

	Mean	Std. Dev.	5th	25th	Median	75th	95th
<i>Panel A: Firm Characteristics</i>							
$XRET$ (%)	0.69	12.31	-17.48	-5.99	-0.16	6.42	21.29
MV (\$ billions)	1.59	5.60	0.01	0.03	0.16	0.81	6.67
BM	0.96	0.82	0.18	0.44	0.75	1.22	2.45
MOM (%)	14.65	49.29	-49.12	-14.24	8.60	34.46	96.30
<i>Panel B: Macroeconomic Variables</i>							
DEF (%)	1.02	0.43	0.55	0.73	0.90	1.21	1.92
DY (%)	3.01	1.10	1.30	2.02	2.96	3.77	4.84
$TBILL$ (%)	5.69	2.70	1.56	4.08	5.16	6.96	10.57
$TERM$ (%)	0.85	1.14	-1.14	0.08	0.78	1.69	2.83

3.5.1 Beta Specification

A key objective in this chapter is to improve the specification of time-varying betas. We investigate whether the time series and cross-sectional variation in betas is best explained by lagged firm characteristics and macroeconomic variables, by past realized betas, or by a linear combination of both. We address this question by estimating the model in equation (3.7) and examining the distribution of ϕ_{it} , which measures the proportion of beta explained by past realized beta. We first compute ϕ_{it} based on equation (3.3) for each draw of the Gibbs sampler. We then calculate for each firm the time series average ϕ_i and its posterior mean.

Figure 3.1 shows the cross-sectional distribution of these posterior means of ϕ_i . The cross-sectional average is 0.51, which implies that for the average firm the mixed beta estimate is the average of the fundamental and realized beta estimates.

The spread in the distribution shows that for some firms past realized betas are more important determinants of mixed betas while for others lagged fundamental betas have a stronger impact.

Figure 3.1: **Cross-Sectional Distribution of Phi**

This figure shows the cross-sectional distribution of the time series average of the parameter ϕ_i , which measures the proportion of beta explained by past realized beta,

$$\beta_{it} = \phi_{it}b_{it} + (1 - \phi_{it})\beta_{it}^*,$$

where b_{it} is the realized beta of firm i , β_{it}^* is the fundamental beta, and where ϕ_{it} is given by

$$\phi_{it} = \phi_{0i} + \phi_{1i}V_{Mt},$$

where V_{Mt} is the realized market variance. We first compute ϕ_{it} for each draw of the Gibbs sampler. We then calculate for each firm the time series average ϕ_i and its posterior mean. This figure shows the cross-sectional distribution of these posterior means.

Figure 3.2 plots the evolution of the cross-sectional average of ϕ_{it} through time. Interestingly, ϕ increases during periods of high market volatility, such as recessions and the crash in 1987. This implies that more weight should be given to past realized betas and less weight to fundamental betas during turbulent conditions. Because the fundamental beta specification is a function of macroeconomic and firm-specific variables, it captures long run movements in beta driven by structural

changes in the economic environment and in firm-specific conditions.⁶ In contrast, because the realized beta specification is based on high-frequency returns, it picks up short run fluctuations in beta in periods of high market volatility.

Figure 3.2: **Evolution of Phi through Time**

This figure plots the evolution through time of the cross-sectional average of ϕ_{it} . We first calculate ϕ_{it} at each iteration of the Gibbs sampler. We then compute its posterior mean and the cross-sectional average of these posterior means in each month from July 1964 through December 2006. Shaded areas indicate NBER recession periods.

Since we find that conditioning variables motivated by economic theory are important determinants of beta, we now consider the posterior distribution of the parameters underlying the fundamental beta. Table 3.2 presents summary statistics of the posterior distribution of the δ_0 and δ_1 parameters in equation (3.6). The constant term δ_0 , which can be interpreted as the average fundamental beta because all conditioning variables are cross-sectionally demeaned, has a posterior mean of 1.01. The posterior distribution of the δ_1 parameters shows that all three firm characteristics are important determinants of fundamental betas. Some of the interaction terms between the firm characteristics and macroeconomic variables also capture important variation in market betas, particularly those involving the default spread and T-Bill rate.

As explained in Section 3.2, we use the MIDAS approach of Ghysels, Santa-

⁶Related to this, Engle and Rangel (2008) model low-frequency patterns in market volatility as a function of macroeconomic and financial variables.

Table 3.2: **Posterior Distribution of Fundamental Beta Parameters**

This table reports the posterior distribution of the determinants of the fundamental beta, which is parameterized as a linear function of firm characteristics and business cycle variables,

$$\beta_{it}^* = \delta_0 + \delta_1' [Z_{it} \otimes BC_t],$$

where Z_{it} is a vector that contains L firm characteristics and BC_t is a vector that contains a constant and M business cycle variables. MV is the log of firm size, BM is the log of the book-to-market ratio, and MOM is the cumulative return over the twelve months prior to the current month. These firm characteristics are expressed as deviations from their cross-sectional mean in every month. DEF is the default spread, DY is the dividend yield, $TBILL$ is the one-month Treasury bill rate, and $TERM$ is the term spread. The table presents the mean, median, standard deviation and 5th, 25th, 75th, and 95th percentile values of the posterior distribution of the δ_0 and δ_1 parameters, based on 5,000 iterations of the Gibbs sampler and a burn-in period of 1,000 iterations. All δ_1 parameters are multiplied by 100 for ease of exposition.

	Mean	Std. Dev.	5th	25th	Median	75th	95th
Constant (δ_0)	1.01	0.01	0.98	0.99	1.00	1.01	1.02
MV	-3.62	0.68	-4.77	-4.05	-3.66	-3.20	-2.47
BM	-5.30	1.32	-7.38	-6.18	-5.44	-4.34	-3.10
MOM	-10.02	2.99	-14.97	-12.05	-9.99	-8.18	-5.28
MV*TBILL	0.65	0.29	0.16	0.45	0.65	0.84	1.14
MV*TERM	1.57	0.28	1.08	1.40	1.58	1.74	2.03
MV*DEF	0.55	0.52	-0.33	0.20	0.53	0.93	1.38
MV*DY	-1.41	0.08	-1.55	-1.47	-1.40	-1.35	-1.27
BM*TBILL	-3.11	0.58	-3.98	-3.52	-3.09	-2.74	-2.20
BM*TERM	0.35	0.73	-0.84	-0.13	0.39	0.83	1.49
BM*DEF	13.15	0.95	11.62	12.53	13.18	13.74	14.64
BM*DY	-0.05	0.14	-0.30	-0.14	-0.04	0.05	0.17
MOM*TBILL	3.05	1.02	1.43	2.29	3.06	3.81	4.67
MOM*TERM	-1.11	0.89	-2.55	-1.67	-1.08	-0.54	0.34
MOM*DEF	-28.91	1.70	-31.65	-30.02	-29.00	-27.76	-26.02
MOM*DY	8.18	0.23	7.81	8.03	8.18	8.32	8.56

Clara, and Valkanov (2005) to estimate realized betas based on daily return data. This approach incorporates a flexible weighting function that makes it possible to choose the optimal weights given to past data in the estimation. The optimal window strikes a balance between giving equal weight to observations to obtain more precise beta estimates and giving more weight to recent data to obtain betas that are timelier and therefore more relevant. As shown in equation (3.5), we use a beta weighting function whose shape is determined by two parameters. We set κ_1 equal to 1 and estimate κ_2 . We find that in our realized beta specification the posterior mean of κ_2 is equal to 1.16. Figure 3.3 compares the optimal weighting scheme implied by this posterior mean of κ_2 to the equal weighting scheme used by traditional rolling window estimators. The plot shows that in the optimal scheme the most recent 150 trading days receive more weight than in the equal weighting scheme because these are most informative for estimating realized betas.

Figure 3.3: Optimal versus Equal Weighting Scheme for Realized Beta

This figure compares the equal weights in the traditional rolling window estimator of realized betas to the estimated optimal weights in the MIDAS estimator of realized betas in equation (3.4). We set the window length equal to 250 trading days.

We now turn to the mixed betas generated by our model. First, we calculate β_{it} based on equation (3.2) at each iteration of the Gibbs sampler. Subsequently, we compute for every firm the time series average of the mixed beta and its posterior mean. Figure 3.4 shows the cross-sectional distribution of these posterior means of β_i . As expected, the distribution is centered around one and has a standard deviation of 0.34. A 95% confidence interval for beta ranges from 0.46 to 1.60, which implies that firms differ substantially in their sensitivity to market movements.

In Table 3.3 we report summary statistics of the posterior means in all three beta specifications. Because the cross-sectional average of beta is close to one in every month, the more interesting aspect is the dispersion of betas, both over time and in the cross-section. The left panel in Table 3.3 reports properties of the cross-section of betas and the right panel shows time series characteristics of beta. The diagonal elements in these panels show that on average, realized betas display the largest spread, both over time and across firms, while fundamental betas show the least variation. This is consistent with the notion that realized

Figure 3.4: **Cross-Sectional Distribution of Firm Betas**

This figure shows the cross-sectional distribution of average firm betas. We first calculate at each iteration of the Gibbs sampler the mixed beta for firm i at time t based on the model in equation (3.7). Subsequently, we compute the time series average of this mixed beta. We then calculate for each firm the posterior mean of its average beta. This figure shows the cross-sectional distribution of these posterior means.

betas capture high-frequency movements in beta whereas fundamental betas pick up long run beta fluctuations. Another explanation is that measurement error in the realized beta estimates leads to spurious variation or that in addition to firm size, book-to-market, and momentum, other firm characteristics drive variation in beta. The time series and cross-sectional behavior of mixed betas is a combination of the dynamics of realized and fundamental betas. Thus, mixed betas combine the benefits of both specifications, responding fast to changes in market conditions without showing excessive variation. The off-diagonal elements in Table 3.3 are the correlations between the betas generated by the three specifications. Fundamental and realized betas are strongly correlated, both over time and across stocks. Correlation is far from perfect though, as a regression of one on the other has an R^2 of only 0.68. This reflects the different cross-sectional characteristics and time series dynamics of realized and fundamental betas. Hence, a combination of these two specifications captures different aspects of market beta dynamics.

Table 3.3: **Beta Summary Statistics**

This table reports summary statistics on the dispersion of betas and on the correlation between mixed, fundamental, and realized betas. The left panel reports the properties of the cross-section of betas based on the time series average of the cross-sectional covariances

$$S_{cross,t} = \frac{1}{N_t} \sum_i \left(\bar{\beta}_{it}^{(j)} - \bar{\bar{\beta}}_t^{(j)} \right) \left(\bar{\beta}_{it}^{(k)} - \bar{\bar{\beta}}_t^{(k)} \right),$$

where the indices j and k refer to the model (Mixed, Realized, Fundamental), betas are evaluated at their posterior means $\bar{\beta}_{it}$, and where $\bar{\bar{\beta}}_t$ is the average beta at time t . The right panel is based on the cross-sectional average of the time series covariances

$$S_{time,i} = \frac{1}{T_i} \sum_t \left(\bar{\beta}_{it}^{(j)} - \bar{\bar{\beta}}_i^{(j)} \right) \left(\bar{\beta}_{it}^{(k)} - \bar{\bar{\beta}}_i^{(k)} \right),$$

where $\bar{\bar{\beta}}_i$ is the average beta of firm i . The diagonal elements in both panels have been transformed into standard deviations. The off-diagonal elements (*in italics*) have been rescaled to correlations.

	Cross-Sectional (S_{cross})			Time Series (S_{time})		
	Mixed	Realized	Fund'l	Mixed	Realized	Fund'l
Mixed beta	0.33			0.24		
Realized beta	<i>0.95</i>	0.61		<i>0.94</i>	0.46	
Fundamental beta	<i>0.94</i>	<i>0.83</i>	0.10	<i>0.93</i>	<i>0.82</i>	0.09

3.5.2 Beta Estimation

In this section we compare the precision of beta estimates from the hierarchical Bayesian panel data model to that of estimates from separate Bayesian time series regressions for all firms. We study the estimation efficiency of the two methods for the fundamental beta specification. Since this specification requires the estimation of many parameters when a large number of conditioning variables is included, we expect the efficiency gain from using the panel model to be substantial. We measure estimation precision by computing the standard deviation and the 5% and 95% percentile values of the posterior distribution of beta at each point in time.

Figure 3.5 plots the posterior mean and 5% and 95% percentile values of the posterior distribution of the fundamental beta of IBM from August 1964 through December 2006. The upper graph is based on the estimation output of the panel data model and the lower graph is constructed using the output of the time series regression. The shaded areas indicate NBER recession periods. The plots show that the confidence interval for beta obtained from the panel regression is much narrower than the interval produced by the time series regression. Noisy estimates of the δ_1 parameters, which measure the influence of conditioning variables on the fundamental beta, lead to wide intervals for beta in the time series model.

The large efficiency gain in the panel model is due to two reasons. First, because the δ coefficients are pooled across stocks in the panel specification, the information in the cross-section of stocks is exploited to obtain more precise estimates. Second,

Figure 3.5: Confidence Interval for IBM Beta: Panel versus Time Series

This figure plots the mean and 5% and 95% percentile values of the posterior distribution of the fundamental beta of IBM in each month from August 1964 through December 2006. The fundamental beta is modeled as a linear function of firm characteristics and macroeconomic variables. The upper graph is based on the estimation output of the hierarchical panel data model presented in Section 3.2 of this chapter and the lower graph is constructed using the output of a time series regression. Shaded areas indicate NBER recession periods.

since we specify hierarchical priors for the firm-specific parameters in the model, only the parameters of the common distribution where the parameters are assumed to be drawn from have to be estimated. The Bayes estimator of the firm-specific parameters in the panel model shrinks the least squares estimator towards the cross-sectional mean. In contrast, in the time series regressions every parameter is estimated individually, which results in poor estimation precision when many parameters need to be estimated and the number of time series observations is small.

Because the panel approach estimates parameters more precisely, it can include more conditioning variables than the traditional approach of estimating a time series regression for every firm used by Avramov and Chordia (2006a). While we include 15 conditioning variables to accurately model beta dynamics, they note that “attention must be restricted to a small number of such variables to ensure some precision in the estimation procedure.” To compare the efficiency of the panel and time series approaches when a more parsimonious specification of fundamental betas is used, we also estimate the panel model and time series regressions with a set of conditioning variables that is similar to that used by Avramov and Chordia (2006a). In particular, we choose firm size, book-to-market, and interactions terms between these characteristics and the default spread as instruments.

The confidence intervals for the fundamental beta of IBM based on this reduced set of conditioning variables are displayed in Figure 3.6. As expected, the intervals for beta generated by the panel model and the time series regression have both narrowed considerably compared to those based on the complete set of conditioning variables. However, the plots show that even when a smaller number of parameters needs to be estimated, the panel approach leads to more precise estimates of firm-specific betas than the time series approach.

Because IBM is present in our data set during the entire sample period, many return observations are available for beta estimation. As explained before, we expect the efficiency gain from the hierarchical Bayesian panel data approach to be even larger for firms with a shorter return history. To summarize the estimation precision of the betas of all firms, we compute the cross-sectional average of the posterior standard deviation of all betas in every month. Figure 3.7 plots these average standard deviations for the panel model and for the time series approach. Clearly, the posterior standard deviation of betas estimated using the time series regressions is larger than the standard deviation of betas estimated using the panel regression.

3.5.3 Portfolio Heterogeneity

The previous section has shown that firm-specific betas are noisy when estimated using time series regressions. To reduce the measurement error in betas, Fama and MacBeth (1973) propose to aggregate stocks into portfolios and run a time series regression for every portfolio to obtain the portfolio’s beta. Fama and French (1992) follow this suggestion and assign each stock the beta of the portfolio it belongs to. As pointed out by Ferson and Harvey (1999), such an approach is

Figure 3.6: **Confidence Interval for Reduced Fundamental Beta of IBM**

This figure plots the mean and 5% and 95% percentile values of the posterior distribution of the fundamental beta of IBM in each month from August 1964 through December 2006. The fundamental beta is modeled as a linear function of a reduced set of firm characteristics and macroeconomic variables. The upper graph is based on the estimation output of the hierarchical panel data model presented in Section 3.2 of this chapter and the lower graph is constructed using the output of a time series regression. Shaded areas indicate NBER recession periods.

Figure 3.7: Average Posterior Standard Deviation of Fundamental Betas

This figure plots the cross-sectional average of the posterior standard deviation of the fundamental betas of all firms in the sample from August 1964 through December 2006. In the upper graph the fundamental beta is modeled as a linear function of firm characteristics and macroeconomic variables and in the lower graph fundamental betas depend on a reduced set of conditioning variables. Posterior standard deviations are based on the estimation output of the hierarchical panel data model and the output of time series regressions estimated for all firms. Shaded areas indicate NBER recession periods.

often used in studies of initial public offerings (IPOs), when no historical return data is available to estimate firm-level betas. They note that this approach only works well when the characteristics used for portfolio formation are good proxies for risk, because an important assumption underlying the portfolio approach is that the stocks in a particular portfolio share the same risk characteristics. In case of the widely used 25 portfolios sorted on firm size and book-to-market, it is assumed that firms are homogeneous in their exposure to risk after controlling for size and B/M. When the stocks in a given portfolio have different exposures to other determinants of risk, this method can lead to serious errors. In this section we therefore examine whether firms that are grouped together in a portfolio have similar risk characteristics.

We construct the 25 size-B/M portfolios following the procedure of Fama and French (1993). Subsequently, we calculate for every portfolio in every month the cross-sectional average and standard deviation of the excess returns and posterior means of the alphas, betas, and phis of the stocks in that portfolio. The left part of Table 3.4 reports for each portfolio the time series means of these cross-sectional averages. Consistent with prior studies (e.g., Fama and French (1996)), the small-growth portfolio has the lowest average return and a large, negative pricing error. In general, average portfolio returns display a strong value premium but weak size effect. Importantly, sorting on firm size and B/M does not produce a wide spread in *average* market betas across portfolios, as most portfolio betas are close to one.⁷ The table further shows that the phi coefficients of large cap portfolios are higher than those of small cap portfolios. This implies that realized betas are more important determinants of the mixed betas of large firms whereas fundamental betas have a stronger effect on the mixed betas of small firms.

Table 3.4 further shows the dispersion of the risk and return characteristics across stocks in each portfolio. For all characteristics we observe strong heterogeneity *within* portfolios. In some portfolios the cross-sectional standard deviation of alphas is more than 1%. Especially firms that are grouped together in small cap portfolios have very different pricing errors. The cross-sectional variation in firm betas within each portfolio is around 0.30, which implies that the assumption that stocks in the same portfolio have similar risk characteristics is violated. Table 3.4 also reports substantial heterogeneity in phi within portfolios. This means that for some firms in a given portfolio mixed betas are mainly driven by realized betas whereas for others fundamental betas are more important.

3.6 Applications of Firm-Specific Betas

This section discusses two important applications of the firm-level betas generated by our Bayesian panel data model. In Section 3.6.1 we compare the explanatory power of different beta specifications and estimation methods for the cross-section

⁷However, unreported results show that value and growth portfolios exhibit different risk *dynamics*. Confirming results of Ang and Chen (2007) and Franzoni (2008), we find that the betas of value firms show a declining trend and are lower than those of growth stocks since the 1980s.

Table 3.4: **Heterogeneity within 25 Size-B/M Portfolios**

This table presents the time series average and cross-sectional spread of characteristics of 25 size-B/M sorted portfolios. We calculate for every portfolio in every month the cross-sectional average and cross-sectional standard deviation of the posterior means of the alphas, betas, and phi and of the excess returns of the stocks in the portfolio. The table shows the time series means of these cross-sectional averages and standard deviations.

Size Quintiles	Book-to-Market Equity (B/M) Quintiles				
	Low	2	3	4	High
Small	Average Return				
	0.18	0.54	0.67	0.97	0.93
	0.72	0.83	0.80	0.83	0.85
	0.67	0.69	0.76	0.95	0.87
	0.65	0.62	0.72	0.84	0.94
Big	Average Return				
	0.48	0.46	0.59	0.72	0.84
	Return Variation				
	15.56	13.84	13.11	12.50	13.83
	12.98	11.16	10.40	10.13	11.34
Small	Alpha Variation				
	1.50	1.27	1.21	1.12	1.28
	1.32	1.07	0.95	0.95	1.13
	1.09	0.88	0.77	0.76	0.91
	0.79	0.61	0.56	0.53	0.59
Big	Alpha Variation				
	0.49	0.40	0.37	0.36	0.40
Small	Beta Variation				
	0.35	0.31	0.31	0.31	0.32
	0.32	0.30	0.29	0.29	0.29
	0.33	0.30	0.29	0.29	0.29
	0.31	0.29	0.29	0.29	0.29
Big	Beta Variation				
	0.27	0.25	0.26	0.25	0.24
Small	Phi Variation				
	0.09	0.10	0.11	0.11	0.11
	0.10	0.11	0.12	0.12	0.11
	0.10	0.11	0.13	0.14	0.12
	0.11	0.12	0.13	0.13	0.11
Big	Phi Variation				
	0.13	0.12	0.11	0.11	0.10

of stock returns. Section 3.6.2 uses the beta estimates to forecast the covariance matrix of returns, which we then use to construct minimum variance portfolios.

3.6.1 Cross-Sectional Tests of the Conditional CAPM

The previous section has shown that aggregating individual stocks into portfolios leads to a substantial loss of information and shrinks the cross-sectional variation in betas. Ang, Liu, and Schwarz (2008) demonstrate that this loss of information can lead to large efficiency losses in cross-sectional tests of asset pricing models. In particular, they show that while creating portfolios reduces estimation error in betas, standard errors of risk premia estimates are higher due to the smaller spread in betas. Consequently, using individual stocks instead of portfolios as base assets allows for more powerful tests of asset pricing models.⁸

Another important reason for using individual stocks in cross-sectional tests of asset pricing models is given by Lewellen, Nagel, and Shanken (2009). They show analytically that due to the strong factor structure in the 25 size-B/M sorted portfolios often used as test assets in asset pricing studies, standard cross-sectional tests have low power to reject a model. In particular, when theoretical restrictions on cross-sectional slopes are ignored, any factor that is only weakly correlated with the true factors can generate a high cross-sectional R^2 and small pricing errors. Lewellen and Nagel (2006) show that the empirical support for several recently proposed asset pricing models weakens considerably when this issue is taken into account. Because individual stock returns do not have a strong factor structure, cross-sectional tests based on individual stocks are not affected by this problem.

In their analysis, Ang, Liu, and Schwarz (2008) assume constant betas, which they estimate by running time series regressions. We extend their work in two directions. First, we improve the specification of betas by allowing for time variation. Second, we use a formal panel data approach to increase the precision of firm-specific beta estimates. We do not claim that the CAPM is superior to competing asset pricing models. Our objective is rather to show the effect of an improved beta specification and estimation on the pricing ability of the CAPM.

We first estimate the betas of all stocks in our sample and of the 25 size-B/M portfolios. We consider four beta specifications (mixed, fundamental, realized, and static) and estimate the models using hierarchical Bayesian panel regressions. These betas are then used as independent variables in second-stage monthly cross-sectional regressions of excess returns on betas,

$$r_{it} = \lambda_{0t} + \lambda_{1t}\beta_{it-1} + \lambda'_{2t}x_{it-1} + v_{it}, \quad (3.10)$$

where λ_{0t} is the intercept, λ_{1t} the risk premium, and where x_{it-1} is a vector of control variables. We run the cross-sectional regressions for every draw of the Gibbs sampler and calculate for each draw the Fama-MacBeth estimator and its

⁸Ang, Liu, and Schwarz (2008) focus on portfolios sorted on firm betas instead of size-B/M sorted portfolios, which should lead to more cross-sectional variation in portfolio betas. However, they find that while ranking stocks on beta produces a large pre-formation spread in beta, the post-formation dispersion is much smaller.

standard error from the time series of the monthly cross-sectional coefficients. We then calculate the posterior mean and variance of the Fama-MacBeth estimator. In Appendix 3.B we demonstrate that this procedure accounts for measurement error in beta by using the entire posterior distribution of the β_{it} in the estimation.

Columns 1-3 in Table 3.5 report the Fama-MacBeth coefficient estimates when individual stocks are used as test assets and no control variables are included in the regression. We find that for the mixed beta specification the intercept is close to zero and insignificant while the risk premium estimate is significantly positive. The λ_1 estimate is 0.56%, which is close to the average monthly excess market return during the sample period (0.47%). This implies that the conditional CAPM with mixed betas satisfies the theoretical restriction emphasized by Lewellen and Nagel (2006) that the risk premium should equal the expected excess factor return. For the other three beta specifications the intercepts are significantly different from zero. The risk premium estimates in the realized beta and fundamental beta models are positive and significant but deviate more from the average market return than the estimated premium in the mixed beta model. The mixed beta specification also outperforms the competing approaches to modeling beta in terms of explanatory power, as measured by the adjusted R^2 . The static CAPM performs worst, because in this model the cross-sectional spread in market betas does not respond to business cycle variations and changes in firm-specific conditions.

Columns 4-6 in Table 3.5 show that when portfolios are used as test assets, all beta specifications generate economically large intercepts. Nevertheless, the mixed beta specification again has the highest explanatory power and the static CAPM does worst. We stress that the R^2 should only be compared across beta specifications and not between individual stocks and portfolios, because the dependent variable in the cross-sectional regressions is different. Table 3.5 further shows that the standard errors of the parameter estimates are much larger when portfolios are used as test assets instead of individual stocks. This result confirms the finding of Ang, Liu, and Schwarz (2008) that sorting stocks into portfolios can lead to large efficiency losses because it reduces the dispersion of betas. In fact, standard errors from using portfolios are more than twice as large as those from using individual stocks.

The last three columns in Table 3.5 report estimation results for individual stocks when control variables are added to the cross-sectional regressions. In particular, the vector x_{it-1} in equation (3.10) contains the firm characteristics size, book-to-market, and momentum. Fama and French (1992) find that the cross-sectional relation between market beta and average return is flat when tests control for size. We find that while adding these firm characteristics leads to an increase in explanatory power, the risk premium estimate for the mixed beta specification remains significantly positive. Thus, when individual stocks are used as test assets and betas are well-specified and precisely estimated, the positive relation between market beta and expected return no longer disappears when controlling for firm characteristics.

3.6.2 Beta Forecasts and Minimum Variance Portfolios

An important application of firm-specific betas is to estimate the covariance matrix of stock returns, which is needed to construct mean-variance efficient portfolios. Traditional implementations of the portfolio theory developed by Markowitz (1952) use sample moments. When the number of assets is large, however, it is difficult to precisely estimate the covariances. As a result, asset weights are often extreme and portfolios behave poorly out-of-sample. Many strategies have been proposed to improve the out-of-sample performance of mean-variance portfolios, such as calculating shrinkage estimators, imposing short selling constraints, and imposing a factor structure on the covariance matrix.⁹

In a recent study, DeMiguel, Garlappi, and Uppal (2009) compare the out-of-sample performance of these approaches to the $1/N$ rule that gives equal weight to all assets. They conclude that none of the more sophisticated methods consistently outperforms the naive $1/N$ benchmark in terms of Sharpe ratio or certainty equivalent return. Of the alternative approaches considered, the global minimum variance portfolio with short selling constraints proposed by Jagannathan and Ma (2003) has the highest Sharpe ratio. The minimum variance portfolio does well because it is the only efficient portfolio that does not require estimates of expected returns, which often contain substantial errors. Chan, Karceski, and Lakonishok (1999) compare the out-of-sample performance of minimum variance portfolios that are based on forecasts of future covariances produced by alternative factor specifications. They find that there is one important factor, the market, that dominates all other factors. Hence, the one-factor model is adequate for forming minimum variance portfolios.

Motivated by these findings, we use the mixed beta estimates produced by the Bayesian panel data model to forecast the covariance matrix of returns and construct the minimum variance portfolio. We expect our method to outperform competing approaches because it delivers more precise estimates of firm-specific betas and because it allows for time variation in these betas. DeMiguel, Garlappi, and Uppal (2009) admit that their assumption of constant risk is a limitation, but argue that models that allow for time-varying moments are likely to perform poorly out-of-sample because many parameters need to be estimated. However, one of the key advantages of our method is that it can estimate a large number of parameters with high precision. We compare the out-of-sample performance of our approach to that of the traditional sample covariance matrix, the static one-factor structure considered by Chan, Karceski, and Lakonishok (1999), and the $1/N$ rule advocated by DeMiguel, Garlappi, and Uppal (2009). Engle and Colacito (2006) stress the importance of isolating the effect of covariance information on portfolio performance from the influence of expected returns if the objective is to evaluate different covariance estimators. Because expected returns do not enter the optimization when constructing minimum variance portfolios, differences in portfolio weights only reflect the effect of differences in covariance forecasts.

⁹See, e.g., Chan, Karceski, and Lakonishok (1999), Jagannathan and Ma (2003), and Ledoit and Wolf (2003).

Table 3.5: Cross-Sectional Asset Pricing Tests

This table reports Fama-MacBeth coefficient estimates for 5,017 NYSE-AMEX stocks (columns 1-3 and 7-9) and for 25 size-B/M portfolios (columns 4-6). The sample period is August 1964 through December 2006. In the first stage, betas are estimated using the Bayesian panel data approach outlined in Section 3.2. These betas are then used as independent variables in second-stage cross-sectional regressions of returns on betas,

$$r_{it} = \lambda_{0t} + \lambda_{1t}\beta_{it-1} + \lambda'_{2t}x_{it-1} + v_{it}.$$

In columns 1-6 no control variables are added. In columns 7-9 the vector x_{it-1} includes the firm characteristics size, book-to-market, and momentum. We use the entire posterior distribution of the betas in the estimation to account for measurement error. The table shows the estimated intercept (λ_0) and risk premium (λ_1) for four different beta specifications. Standard errors are in parentheses and corresponding t -statistics are in brackets.

	Stocks - No Controls			Portfolios - No Controls			Stocks - With Controls		
	λ_0	λ_1	Adj. R^2	λ_0	λ_1	Adj. R^2	λ_0	λ_1	Adj. R^2
<i>Panel A: Mixed Beta</i>									
Estimate	-0.04	0.56	4.21	0.61	0.08	18.30	0.04	0.63	6.70
SE	(0.17)	(0.26)		(0.42)	(0.48)		(0.17)	(0.30)	
t -stat	[-0.24]	[2.15]		[1.45]	[0.16]		[0.22]	[2.09]	
<i>Panel B: Realized Beta</i>									
Estimate	0.35	0.35	3.47	0.61	0.13	16.66	0.31	0.41	6.26
SE	(0.13)	(0.16)		(0.29)	(0.29)		(0.12)	(0.16)	
t -stat	[2.73]	[2.24]		[2.07]	[0.44]		[2.51]	[2.48]	
<i>Panel C: Fundamental Beta</i>									
Estimate	0.64	0.67	2.05	1.14	-0.42	16.24	0.72	0.72	4.55
SE	(0.30)	(0.34)		(0.76)	(0.78)		(0.24)	(0.40)	
t -stat	[2.13]	[1.97]		[1.50]	[-0.54]		[3.04]	[1.80]	
<i>Panel D: Static Beta</i>									
Estimate	0.63	0.09	0.65	1.22	-0.45	9.39	0.87	-0.15	3.78
SE	(0.28)	(0.20)		(0.47)	(0.49)		(0.29)	(0.22)	
t -stat	[2.25]	[0.45]		[2.60]	[-0.92]		[2.98]	[-0.70]	

The first estimator of the covariance matrix is the sample covariance matrix,

$$S_t^S = \frac{1}{T-1} \sum_{t=1}^T (R_t - \bar{R})(R_t - \bar{R}), \quad (3.11)$$

where R_t is the vector of monthly stock returns and $\bar{R}$ contains the sample mean returns. The second covariance estimator is based on the one-factor model. In the first step, we estimate mixed betas using our panel approach and static betas using time series regressions. We use these betas to estimate the covariance matrix according to the one-factor model,

$$S_t^F = s_{Mt}^2 B_t B_t' + D, \quad (3.12)$$

where B_t is the $N_t \times 1$ vector of betas, s_{Mt}^2 is the sample variance of the market return, and D is a diagonal matrix that contains the variances of the residuals.¹⁰

Subsequently, we use the various estimates of the covariance matrix to construct the minimum variance portfolio, by choosing the portfolio weights that solve the following problem

$$\min w_t' S_t w_t, \quad (3.13)$$

$$\text{s.t. } \sum_i w_{it} = 1. \quad (3.14)$$

The constraint ensures that the portfolio is fully invested. Following Chan, Karceski, and Lakonishok (1999) and Jagannathan and Ma (2003), we also consider an extension in which we add a short selling constraint,

$$w_{it} \geq 0, \quad i = 1, 2, \dots, N. \quad (3.15)$$

We form the minimum variance portfolio at the end of each month based on the forecast of the covariance matrix for the next month. The first portfolio is formed using the first half of the sample period to estimate the covariance matrix of returns according to the methods described above. Because the sample covariance matrix cannot be positive definite unless the number of return observations per stock is larger than the number of stocks, we only apply the sample covariance approach to a subset of the investment universe. We record the performance of the minimum variance portfolio in the next month and rebalance the portfolio using the new forecast of the covariance matrix. This procedure generates a time series of monthly returns for global minimum variance portfolios constructed using different covariance estimators. As a benchmark we also form an equal-weighted portfolio at the end of each month. Since the objective is to minimize the portfolio variance, we evaluate the performance of the different methods by calculating the realized volatility of the portfolio returns.

¹⁰The assumption that the residuals are cross-sectionally uncorrelated follows Jagannathan and Ma (2003) and boils down to assuming that the one-factor model captures the cross-section of returns. For simplicity we stick to the one-factor structure but a natural extension would be to consider multifactor models to pick up any remaining cross-sectional dependence in the residuals.

Table 3.6 reports annualized risk and return characteristics of the minimum variance portfolios constructed using various forecasts of the covariance matrix. Panel A shows the out-of-sample performance when all stocks in the sample are used in the optimization and short selling is allowed. The mixed beta specification estimated using the panel data approach outperforms all other methods and produces a portfolio with an annualized standard deviation of 8.12%. The $1/N$ strategy leads to a standard deviation that is almost twice as large (15.40%). The static beta model ranks second and generates an out-of-sample standard deviation of 8.50%. Although our sole objective is to minimize portfolio variance without taking portfolio return into consideration, we find that the mixed beta approach also leads to the best risk-return tradeoff, as it produces the highest Sharpe ratio (0.69). Moreover, since the minimum variance portfolio constructed using the mixed beta approach does not involve short positions, it can be easily implemented in practice.

In Panel B we report the performance for portfolios constructed from a random sample of 250 stocks, which is the same number of stocks considered by Chan, Karceski, and Lakonishok (1999). For this subset of stocks the $1/N$ rule leads to the highest out-of-sample standard deviation (15.93%), followed by the sample covariance matrix, which generates a standard deviation of 15.19% and takes large short positions, and the static one-factor model that produces a portfolio with a standard deviation equal to 13.32%. The mixed beta approach again beats all other methods and yields a portfolio with a standard deviation of 8.41%. Thus, its performance when applied to the smaller sample of stocks is similar to that when all stocks are used in the optimization. An important reason for the relatively poor performance of the naive $1/N$ strategy is that we allocate wealth across individual stocks. In contrast, DeMiguel, Garlappi, and Uppal (2009) apply this policy to allocate wealth across portfolios of stocks. They point out that the loss from naive as opposed to optimal diversification is much larger when allocating wealth across individual assets, because individual stocks have higher idiosyncratic volatility than portfolios.

Jagannathan and Ma (2003) document that the out-of-sample performance of the sample covariance approach can be improved by imposing no-short-sale constraints, because these limit the effect of sampling error on portfolio weights. Panel C reports the risk and return characteristics of the portfolios generated by the four methods when the nonnegativity constraint is imposed on the weights. The random sample of 250 stocks used to form the portfolios is the same as that in Panel B. We find that the standard deviation of the minimum variance portfolio constructed using the sample covariance matrix is indeed lower when no-short-sale restrictions are in place. In fact, this method yields a lower standard deviation than the equal-weighted portfolio or the portfolio based on the one-factor model with static betas. We confirm the finding of Jagannathan and Ma (2003) that imposing no-short-sale constraints reduces the performance of the static factor model. Because the portfolio produced by the mixed beta approach in Panel B does not take short positions, it is not affected by the nonnegativity constraint and still has the smallest out-of-sample standard deviation.

Table 3.6: **Risk and Return of Minimum Variance Portfolios**

This table reports the out-of-sample performance of global minimum variance portfolios that are formed at the end of each month from December 1985 through December 2006 out of a universe of NYSE-AMEX stocks. The optimization procedure uses forecasts of the covariance matrix of returns produced by different methods. Panel A reports the out-of-sample performance of minimum variance portfolios when all stocks in the sample are used in the optimization and without any constraints imposed on the weights. Panel B reports results for the unconstrained optimization when a random sample of 250 stocks is used to construct the portfolios. Panel C reports results for this reduced investment universe when a nonnegativity (no-short-sale) constraint is imposed on the portfolio weights. Mean excess returns, standard deviations, and Sharpe ratios are annualized. Short interest is in percentages.

	Mean	Std. Dev.	Sharpe Ratio	Short Interest
<i>Panel A: Unconstrained (all stocks)</i>				
Equal-weighted ($1/N$)	8.57	15.40	0.56	0.00
Static beta (TS model)	5.24	8.50	0.61	-64.64
Mixed beta (Panel model)	5.60	8.12	0.69	0.00
<i>Panel B: Unconstrained (250 stocks)</i>				
Sample covariance matrix	4.30	15.19	0.28	-144.79
Equal-weighted ($1/N$)	9.46	15.93	0.59	0.00
Static beta (TS model)	7.81	13.32	0.59	-60.35
Mixed beta (Panel model)	5.70	8.41	0.68	0.00
<i>Panel C: Nonnegativity Constrained (250 stocks)</i>				
Sample covariance matrix	3.02	11.74	0.26	0.00
Equal-weighted ($1/N$)	9.46	15.93	0.59	0.00
Static beta (TS model)	7.55	15.60	0.48	0.00
Mixed beta (Panel model)	5.70	8.41	0.68	0.00

3.7 Conclusion

Many applications of modern finance theory require precise beta estimates of individual stocks. However, as noted by Campbell, Lettau, Malkiel, and Xu (2001), “firm-specific betas are difficult to estimate and may well be unstable over time.” Academics and practitioners have taken two approaches to estimating firm betas. The first method sorts stocks into portfolios based on firm characteristics to reduce measurement error and assigns each firm the beta of the portfolio it belongs to. However, when stocks in the same portfolio have different exposures to other determinants of risk than the characteristics they are sorted on, this approach can lead to serious errors. The second method estimates a separate time series regression for every stock. Although this approach allows each firm to have its own risk exposure, the resulting estimates can be very noisy. The literature also uses different methods to model time variation in betas. Many studies use a parametric

approach in which variation in beta is modeled as a linear function of conditioning variables. An alternative, nonparametric approach to model risk dynamics is based on data-driven filters. However, both methods have important drawbacks and involve a tradeoff between the precision and timeliness of beta estimates.

In this chapter we therefore improve both the specification and estimation of time-varying, firm-specific betas. We combine the parametric and nonparametric approaches for modeling changes in betas. The precision of firm-level beta estimates is increased by setting up a Bayesian panel data model that exploits the information contained in the cross-section of stocks and imposes a common structure on parameters while still allowing for cross-sectional heterogeneity.

We find that modeling time-varying betas as a function of both conditioning variables and past returns is preferred over traditional specifications that are based on only one of these components. Because fundamental and realized betas exhibit different time series dynamics and cross-sectional characteristics, a combination of these specifications captures different aspects of beta. We show that the optimal mixture of the two betas varies across firms and over time. We further demonstrate that our panel approach yields more precise estimates of stock betas than the traditional approach of estimating betas by running a time series regression for every firm. Moreover, we document strong cross-sectional variation in betas of firms that are grouped together in portfolios sorted on size and book-to-market. Consequently, aggregating stocks into portfolios conceals important information contained in individual betas and reduces the cross-sectional dispersion of betas.

We demonstrate that the mixed betas generated by our panel data model lead to a sharp increase in the pricing ability of the conditional CAPM. The estimate of the market risk premium remains significantly positive when controlling for firm characteristics. The results also support the finding of Ang, Liu, and Schwarz (2008) that the use of individual stocks as tests assets instead of portfolios leads to more efficient estimates in cross-sectional tests of asset pricing models. We extend their work by showing that a better specification and more precise estimation of stock-specific betas increases the explanatory power of the CAPM.

Accurate estimates of firm-specific betas are also important for portfolio optimization. Based on the mixed beta estimates produced by our panel model we forecast the covariance matrix of stock returns, which is then used to form minimum variance portfolios. The portfolio constructed using the mixed betas outperforms portfolios based on other strategies, such as the traditional sample covariance matrix and the naive $1/N$ rule, in terms of out-of-sample standard deviation, even after imposing short selling constraints.

Since our framework is flexible, it can be readily extended to include multiple risk factors, a different set of conditioning variables for fundamental betas, or another window length for estimating realized betas. In addition, while we have demonstrated the advantages of our approach for asset pricing and portfolio management, it also has important benefits for corporate finance applications. For example, because it quickly captures changes in beta and generates precise beta estimates even when little return data is available, our method is well suited for calculating risk-adjusted returns in studies of IPOs and M&As.

3.A Appendix A: Posterior Distributions

3.A.1 Joint Posterior Distribution

The joint posterior density is proportional to the product of the likelihood function and the prior distributions of all parameters in the set θ , $p(\theta|y) \propto p(y|\theta)p(\theta)$. Defining β_{it} as in equation (3.2), stacking the time series observations for every firm i into vectors, and substituting the prior densities specified in Section 3.3.2, produces the following joint posterior distribution¹¹

$$\begin{aligned}
 p(\theta|y) &= p(\alpha_i, \sigma_\alpha^2, \phi_{0i}, \sigma_{\phi_0}^2, \phi_1, \sigma_{\phi_1}^2, \kappa_2, \delta_0, \sigma_{\delta_0}^2, \delta_1, \Omega_{\delta_1}^{-1}, \sigma_{\epsilon_i}^2 | y) \\
 &\propto \prod_{i=1}^N (\sigma_{\epsilon_i}^2)^{-\frac{T_i}{2}} \exp \left[-\frac{1}{2\sigma_{\epsilon_i}^2} (r_i - \alpha_i - r_M \beta_i)' (r_i - \alpha_i - r_M \beta_i) \right] \\
 &\times \prod_{i=1}^N (\sigma_\alpha^2)^{-\frac{1}{2}} \exp \left[-\frac{\alpha_i^2}{2\sigma_\alpha^2} \right] \times (\sigma_\alpha^2)^{A_\alpha+1} \exp [-\sigma_\alpha^{-2} B_\alpha] \\
 &\times \prod_{i=1}^N (\sigma_{\phi_0}^2)^{-\frac{1}{2}} \exp \left[-\frac{(\phi_{0i} - \mu_{\phi_0})^2}{2\sigma_{\phi_0}^2} \right] \times (\sigma_{\phi_0}^2)^{A_{\phi_0}+1} \exp [-\sigma_{\phi_0}^{-2} B_{\phi_0}] \\
 &\times (\sigma_{\phi_1}^2)^{-\frac{1}{2}} \exp \left[-\frac{\phi_1^2}{2\sigma_{\phi_1}^2} \right] \times (\sigma_{\phi_1}^2)^{A_{\phi_1}+1} \exp [-\sigma_{\phi_1}^{-2} B_{\phi_1}] \\
 &\times (\sigma_{\delta_0}^2)^{-\frac{1}{2}} \exp \left[-\frac{\delta_0^2}{2\sigma_{\delta_0}^2} \right] \times (\sigma_{\delta_0}^2)^{A_{\delta_0}+1} \exp [-\sigma_{\delta_0}^{-2} B_{\delta_0}] \\
 &\times |\Omega_{\delta_1}^{-1}|^{\frac{1}{2}} \exp \left[-\frac{1}{2} \delta_1' \Omega_{\delta_1}^{-1} \delta_1 \right] \times |\Omega_{\delta_1}^{-1}|^{\frac{\psi_{\delta_1} - (L+LM) - 1}{2}} \exp \left[-\frac{1}{2} \text{tr} ([\psi_{\delta_1} S_{\delta_1}] \Omega_{\delta_1}^{-1}) \right] \\
 &\times \prod_{i=1}^N (\sigma_{\epsilon_i}^2)^{A_\epsilon+1} \exp [-\sigma_{\epsilon_i}^{-2} B_{\epsilon_i}].
 \end{aligned}$$

¹¹We use the following parameterization of the inverse gamma distribution,

$$p(y|A, B) = \frac{B^A}{\Gamma(A)} \left(\frac{1}{y^{A+1}} \right)^{A+1} \exp \left(\frac{-B}{y} \right),$$

where $\Gamma(A)$ denotes the Gamma function, A is the shape parameter, and B is the scale parameter. For the Wishart distribution we use the parameterization,

$$p(H|R, \nu) \propto \frac{|H|^{(\nu-k-1)/2}}{|R|^{\nu/2}} \exp \left[\frac{1}{2} \text{tr} (R^{-1} H) \right],$$

where k denotes the dimension of the matrix H , ν is the degrees of freedom parameter, and R is the scale matrix.

3.A.2 Conditional Posterior Distributions

In order to implement the Gibbs sampler we need to derive the full conditional posterior densities for each block of parameters. The conditional densities can be derived from the joint posterior density by ignoring all terms that do not depend on the parameters of interest and treating the parameters considered to be known as constants. We then obtain the conditional density for the parameters of interest by rearranging the remaining terms into the kernel of a known distribution. We partition the set of all parameters θ into the following blocks

θ_1 : MIDAS weight parameter: κ_2

θ_2 : Alpha parameters: $\{\alpha_i\}$

θ_3 : Fundamental beta parameters: δ_0, δ_1

θ_4 : Firm-specific mixed beta parameters: $\{\phi_{0i}\}$

θ_5 : Pooled mixed beta parameter: ϕ_1

θ_6 : Variance and covariance parameters: $\sigma_\alpha^2, \sigma_{\delta_0}^2, \Omega_{\delta_1}^{-1}, \sigma_{\phi_0}^2, \sigma_{\phi_1}^2, \sigma_{\epsilon_i}^2$

To generate samples from the conditional posterior of θ_1 we use the Metropolis-Hastings algorithm. The conditional posteriors for all other blocks have convenient functional forms. Therefore, we can use the Gibbs sampler to iteratively draw from the conditional densities of $\theta_2, \theta_3, \theta_4, \theta_5$, and θ_6 . In subsequent derivations we write θ_{-x} to denote all elements of θ other than x . To simplify notation, we rewrite the model in equation (3.7) in matrix form as

$$r_i = \alpha_i \iota_{T_i} + \phi_{0i} r_M b_i + \phi_1 r_M V_M b_i + r_M W_i \delta - \phi_{0i} r_M W_i \delta - \phi_1 r_M V_M W_i \delta + \epsilon_i, \quad (3.A.1)$$

where r_i is a $T \times 1$ vector of excess returns, r_M a $T \times T$ diagonal matrix of excess market returns, V_M a $T \times T$ diagonal matrix of lagged market volatility, b_i a $T \times 1$ vector of realized betas, and ϵ_i a $T \times 1$ vector of idiosyncratic shocks. Since the δ_0 and δ_1 parameters in block θ_3 have independent priors, we have simplified the notation further by rewriting $\delta_0 \iota_T + ZBC_i \delta_1$ as $W_i \delta$, where W_i is the $T \times (1 + L + LM)$ matrix of the constant term and conditioning variables. We combine the corresponding precisions $\sigma_{\delta_0}^{-2}$ and $\Omega_{\delta_1}^{-1}$ into the matrix Ω_δ^{-1} .

Metropolis-Hastings algorithm to draw κ_2

Since we implement a change of variable, $\kappa_2 = 1 + 25\tilde{\kappa}_2$, we need to draw values for $\tilde{\kappa}_2$. Because the conditional posterior density for $\tilde{\kappa}_2$ does not take a standard form, we cannot use the Gibbs sampler. Instead, we employ the Metropolis-Hastings (M-H) algorithm, which is a general accept-reject algorithm. In fact, Gelman, Carlin, Stern, and Rubin (2004) show that the Gibbs sampler is a special case of Metropolis-Hastings in which proposed parameter values are accepted with probability one. The M-H algorithm proceeds as follows.

First, a candidate value $\tilde{\kappa}_2^*$ is drawn from a proposal density $q(\tilde{\kappa}_2)$. We apply the Independence Chain M-H algorithm, in which the proposal density is independent across draws. We choose a Beta(1,3) proposal density, which has a mean of 0.25 and standard deviation equal to 0.19. Because the proposal density is not

identical to the posterior density, the M-H algorithm does not accept all proposal draws. When a proposal is rejected the parameter value is set equal to the current value. Draws are accepted according to the following probability

$$\pi(\tilde{\kappa}_2^{(g-1)}, \tilde{\kappa}_2^*) = \min \left\{ 1, \frac{p(\tilde{\kappa}_2^*|y)q(\tilde{\kappa}_2^{(g-1)})}{p(\tilde{\kappa}_2^{(g-1)}|y)q(\tilde{\kappa}_2^*)} \right\}, \quad (3.A.2)$$

where $\tilde{\kappa}_2^{(g-1)}$ is the current parameter value. This approach ensures that candidate draws with a high posterior density have a higher probability of being accepted than draws with a low posterior density. Repeating this procedure produces the required sequence of draws from the posterior distribution.

Conditional posterior α_i

Using Bayes' theorem, we can write:

$$\begin{aligned} p(\alpha_i | \theta_{-\alpha_i}, y) &\propto p(y | \alpha_i, \theta_{-\alpha_i}) p(\alpha_i) \\ &\propto \exp \left[-\frac{1}{2} Q^* \right], \end{aligned}$$

$$\begin{aligned} \text{where } Q^* &= (r_i - \alpha_i - r_M \beta_i)' \Omega_{\epsilon_i}^{-1} (r_i - \alpha_i - r_M \beta_i) + \frac{\alpha_i^2}{\sigma_\alpha^2} \\ &= (X_{\alpha_i} - \alpha_i \iota_{T_i})' \Omega_{\epsilon_i}^{-1} (X_{\alpha_i} - \alpha_i \iota_{T_i}) + \frac{\alpha_i^2}{\sigma_\alpha^2} \\ &= Q_1^* + Q_2^*, \end{aligned}$$

$$\text{with } Q_1^* = \frac{(\alpha_i - \bar{\alpha}_i)^2}{\bar{\sigma}_{\alpha_i}^2},$$

$$\text{and } Q_2^* = X_{\alpha_i}' \Omega_{\epsilon_i}^{-1} X_{\alpha_i} - \frac{\bar{\alpha}_i^2}{\bar{\sigma}_{\alpha_i}^2},$$

and where $X_{\alpha_i} = r_i - \phi_{0i} r_M b_i - \phi_{1i} r_M V_M b_i - r_M W_i \delta + \phi_{0i} r_M W_i \delta + \phi_{1i} r_M V_M W_i \delta$.

In the derivation of $p(\alpha_i | \theta_{-\alpha_i}, y)$ all parameters in Q_2^* are known, so we can treat Q_2^* as a constant. Thus, $p(\alpha_i | \theta_{-\alpha_i}, y)$ is proportional to $\exp[-\frac{1}{2} Q_1^*]$, which is the kernel of a normal density. Therefore,

$$\begin{aligned} \alpha_i | \theta_{-\alpha_i}, y &\sim N(\bar{\alpha}_i, \bar{\sigma}_{\alpha_i}^2), \\ \text{with } \bar{\alpha}_i &= \left[\iota_{T_i}' \Omega_{\epsilon_i}^{-1} \iota_{T_i} + \frac{1}{\sigma_\alpha^2} \right]^{-1} [\iota_{T_i}' \Omega_{\epsilon_i}^{-1} X_{\alpha_i}], \\ \text{and } \bar{\sigma}_{\alpha_i}^2 &= \left[\iota_{T_i}' \Omega_{\epsilon_i}^{-1} \iota_{T_i} + \frac{1}{\sigma_\alpha^2} \right]^{-1}. \end{aligned}$$

Conditional posterior δ

$$\begin{aligned}
 \delta | \theta_{-\delta}, y &\sim N(\bar{\delta}, \bar{\Omega}_\delta), \\
 \text{with } \bar{\delta} &= \left[\sum_{i=1}^N ((1 - \phi_{0i}) r_M W_i - \phi_1 r_M V_M W_i)' \Omega_{\epsilon_i}^{-1} \right. \\
 &\quad \times ((1 - \phi_{0i}) r_M W_i - \phi_1 r_M V_M W_i) + \Omega_\delta^{-1} \left. \right]^{-1} \\
 &\quad \times \left[\sum_{i=1}^N X'_{\delta i} \Omega_{\epsilon_i}^{-1} ((1 - \phi_{0i}) r_M W_i - \phi_1 r_M V_M W_i) \right], \\
 \text{and } \bar{\Omega}_\delta &= \left[\sum_{i=1}^N ((1 - \phi_{0i}) r_M W_i - \phi_1 r_M V_M W_i)' \Omega_{\epsilon_i}^{-1} \right. \\
 &\quad \times ((1 - \phi_{0i}) r_M W_i - \phi_1 r_M V_M W_i) + \Omega_\delta^{-1} \left. \right]^{-1},
 \end{aligned}$$

and where $X_{\delta i} = r_i - \alpha_i \iota_{T_i} - \phi_{0i} r_M b_i - \phi_1 r_M V_M b_i$.

Conditional posterior ϕ_{0i}

$$\begin{aligned}
 \phi_{0i} | \theta_{-\phi_{0i}}, y &\sim N(\bar{\phi}_{0i}, \bar{\sigma}_{\phi_{0i}}^2), \\
 \text{with } \bar{\phi}_{0i} &= \left[(b_i - W_i \delta)' r_M \Omega_{\epsilon_i}^{-1} r_M (b_i - W_i \delta) + \frac{1}{\sigma_{\phi_0}^2} \right]^{-1} \\
 &\quad \times \left[(b_i - W_i \delta)' r_M \Omega_{\epsilon_i}^{-1} X_{\phi_{0i}} + \frac{\mu_{\phi_0}}{\sigma_{\phi_0}^2} \right], \\
 \text{and } \bar{\sigma}_{\phi_{0i}}^2 &= \left[(b_i - W_i \delta)' r_M \Omega_{\epsilon_i}^{-1} r_M (b_i - W_i \delta) + \frac{1}{\sigma_{\phi_0}^2} \right]^{-1}, \\
 \text{and where } X_{\phi_{0i}} &= r_i - \alpha_i \iota_{T_i} - r_M W_i \delta - \phi_1 r_M V_M (b_i - W_i \delta).
 \end{aligned}$$

Conditional posterior ϕ_1

$$\phi_1 | \theta_{-\phi_1}, y \sim N(\bar{\phi}_1, \bar{\sigma}_{\phi_1}^2),$$

$$\text{with } \bar{\phi}_1 = \left[\sum_{i=1}^N (b_i - W_i \delta)' r_M V_M \Omega_{\epsilon_i}^{-1} r_M V_M (b_i - W_i \delta) + \frac{1}{\sigma_{\phi_1}^2} \right]^{-1} \\ \times \left[\sum_{i=1}^N X'_{\phi_{1i}} \Omega_{\epsilon_i}^{-1} r_M V_M (b_i - W_i \delta) \right],$$

$$\text{and } \bar{\sigma}_{\phi_1}^2 = \left[\sum_{i=1}^N (b_i - W_i \delta)' r_M V_M \Omega_{\epsilon_i}^{-1} r_M V_M (b_i - W_i \delta) + \frac{1}{\sigma_{\phi_1}^2} \right]^{-1},$$

$$\text{and where } X_{\phi_{1i}} = r_i - \alpha_i \iota_{T_i} - \phi_{0i} r_M (b_i - W_i \delta) - r_M W_i \delta.$$

Conditional posteriors $\sigma_\alpha^2, \Omega_\delta^{-1}, \sigma_{\phi_0}^2, \sigma_{\phi_1}^2, \sigma_{\epsilon_i}^2$

$$\sigma_\alpha^2 | \theta_{-\theta_6}, y \sim IG \left(\frac{N + 2A_\alpha}{2}, \frac{\sum_{i=1}^N \alpha_i^2 + 2B_\alpha}{2} \right),$$

$$\Omega_\delta^{-1} | \theta_{-\theta_6}, y \sim Wish \left([\delta \delta' + (\psi_\delta S_\delta)]^{-1}, \psi_\delta + 1 \right),$$

$$\sigma_{\phi_0}^2 | \theta_{-\theta_6}, y \sim IG \left(\frac{N + 2A_{\phi_0}}{2}, \frac{\sum_{i=1}^N (\phi_{0i} - \mu_{\phi_0})^2 + 2B_{\phi_0}}{2} \right),$$

$$\sigma_{\phi_1}^2 | \theta_{-\theta_6}, y \sim IG \left(\frac{1 + 2A_{\phi_1}}{2}, \frac{\phi_1^2 + 2B_{\phi_1}}{2} \right),$$

$$\sigma_{\epsilon_i}^2 | \theta_{-\theta_6}, y \sim IG \left(\frac{T_i + 2A_\epsilon}{2}, \frac{(r_i - \alpha_i - r_M \beta_i)' (r_i - \alpha_i - r_M \beta_i) + 2B_\epsilon}{2} \right).$$

3.B Appendix B: Cross-Sectional Tests

In this appendix we show how we account for measurement error in betas in the cross-sectional asset pricing tests. We consider the cross-sectional regression model described in Section 3.6,

$$r_{it} = \lambda_{0t} + \lambda_{1t}\beta_{it-1} + \lambda'_{2t}x_{it-1} + v_{it}. \quad (3.B.1)$$

Conditional on β_{it-1} we can estimate the parameters λ_t , either Bayesian or classical, using the Fama-MacBeth approach. Let $W_{it} = (1 \ \beta_{it} \ x'_{it})'$ and let $\hat{\lambda}_t$ denote the cross-sectional OLS estimator of λ_t . The Fama-MacBeth estimator of λ is

$$\hat{\lambda} = \frac{1}{T} \sum_t \hat{\lambda}_t = \frac{1}{T} (W'_{t-1} W_{t-1})^{-1} W'_{t-1} r_t, \quad (3.B.2)$$

with covariance matrix

$$S \equiv Var \left(\sqrt{T}(\hat{\lambda} - \lambda) \right) = \frac{1}{T} \sum_t (\hat{\lambda}_t - \hat{\lambda})(\hat{\lambda}_t - \hat{\lambda})'. \quad (3.B.3)$$

The Fama-MacBeth procedure overstates the precision of parameter estimates in the cross-sectional regressions because it ignores estimation errors in the β_{it} .

As explained in Section 3.3, the Gibbs sampler produces a series of L draws from the posterior density $p(\beta|y)$, where β contains the entire collection of all β_{it} and y is a shorthand for all data used in estimating the betas. Given the $\beta_{it}^{(\ell)}$ from the ℓ^{th} iteration of the Gibbs sampler we can form the regressor matrix $W_t^{(\ell)}$ and use $W_t^{(\ell)}$ to construct the conditional mean $\hat{\lambda}^{(\ell)}$ and covariance matrix $S^{(\ell)}$. From these we form the unconditional estimators

$$\bar{\lambda} = \frac{1}{L} \sum_{\ell} \hat{\lambda}^{(\ell)} \quad (3.B.4)$$

and

$$\bar{S} = \frac{1}{L} \sum_{\ell} S^{(\ell)} + \frac{1}{L} \sum_{\ell} (\hat{\lambda}^{(\ell)} - \bar{\lambda})(\hat{\lambda}^{(\ell)} - \bar{\lambda})'. \quad (3.B.5)$$

The estimates $\bar{\lambda}$ and $\bar{S}$ can be interpreted as the posterior mean and variance of λ if we assume that the prior on λ_t is uniform and that λ_t does not affect the posterior density of β_{it} .

Chapter 4

Long and Short Run Correlation Risk in Stock Returns

In this chapter I study the pricing of long and short run variance and correlation risk in the time series of aggregate stock returns and in the cross-section of individual returns. I find that the predictive power of the market variance risk premium for the equity premium is completely driven by the correlation risk premium. Furthermore, I show that long run market volatility risk is priced in the cross-section because of priced innovations in long-term idiosyncratic volatility and shocks to long-term market-wide correlations. In contrast, short run market volatility risk is only priced because of short-term correlation risk.

4.1 Introduction

Traditionally, risk in financial markets is measured as the risk that returns vary over time. However, a large body of empirical evidence shows that risk itself also changes over time. During the credit crisis of 2008, for example, implied market volatility increased from 20% at the end of August to 80% two months later. Since market volatility depends on the correlations between all stocks in the market and the volatility of these stocks, changes in market risk reflect changes in market-wide correlations, changes in individual risk, or changes in both.¹ Because time-varying market risk affects the future risk-return tradeoff, it can be a priced risk factor in intertemporal models. The decomposition of market risk implies that if market variance risk is priced, individual variance risk, correlation risk, or both should be

¹Surveys of the vast literature on time-varying volatility are given by Bollerslev, Engle, and Nelson (1994) and Ghysels, Harvey, and Renault (1996). Empirical evidence that correlations vary over time and tend to increase when stock prices decrease is presented in Bollerslev, Engle, and Wooldridge (1988) and Moskowitz (2003).

priced. Intuitively, investors prefer stocks that pay off when aggregate uncertainty is high and when diversification opportunities are reduced by a sudden increase in correlations. Furthermore, because different determinants of risk induce high- and low-frequency movements in volatilities and correlations, variance and correlation risk can be priced at multiple horizons.²

In this chapter I analyze the pricing of long and short run variance and correlation risk in aggregate returns and in the cross-section of individual stock returns. At the market level, Bollerslev, Tauchen, and Zhou (2009) show that the variance risk premium, measured as the difference between implied and realized market variance, predicts the equity premium. However, recent work by Driessen, Maenhout, and Vilkov (2009) shows that market variance risk is only priced in option markets because of priced correlation risk, not priced individual variance risk. Motivated by these findings, my first objective in this chapter is to study whether the correlation risk premium and individual variance risk premium predict stock market returns. Using data for S&P 100 index options and for individual options on the S&P 100 components, I show that the predictive power of the market variance premium completely derives from the predictive power of the correlation risk premium. In fact, for one-month ahead returns, the magnitude of the predictive power of the correlation premium is larger than that of the market variance premium. The predictive power is even stronger at a quarterly horizon due to the persistence of the correlation risk premium and is robust to the inclusion of traditional return predictors like the default spread, price-earnings ratio, consumption-wealth ratio, and term spread. In contrast to the strong predictive power of the correlation risk premium, the average individual variance risk premium does not predict the equity premium at the monthly or quarterly horizon.

My second contribution in this chapter is to examine whether long and short run components of correlation risk and idiosyncratic volatility risk are priced in the cross-section of returns. The finding that at the aggregate level, variance risk is priced because of correlation risk suggests that the cross-sectional relation between innovations in market volatility and returns documented by Ang, Hodrick, Xing, and Zhang (2006) is also driven by innovations in market-wide correlations. In that case, stocks that pay off when correlations are high should have lower expected returns. However, since market volatility is also determined by individual volatilities, the *innovation* in average idiosyncratic risk can be a priced risk factor in the cross-section of returns as well.³ Moreover, the finding of Adrian and Rosenberg (2008) that shocks to long and short run components of market volatility carry a different price of risk raises the question whether correlation risk and idiosyncratic risk are also priced at different horizons.

²Engle and Lee (1999) and Chernov, Gallant, Ghysels, and Tauchen (2003) propose long and short run component models of volatility. Campbell, Lettau, Malkiel, and Xu (2001) find evidence of an upward trend in idiosyncratic volatility and a downward trend in correlations until 1997. Engle and Rangel (2008) link variation in long run market volatility to macroeconomic variables.

³This differs from analyzing the relation between the *level* of idiosyncratic risk of a particular firm and its return, as in Ang, Hodrick, Xing, and Zhang (2006) and Fu (2009). These papers are motivated by the incomplete information model of Merton (1987), in which idiosyncratic risk is priced because investors do not hold diversified portfolios.

To study these cross-sectional predictions, I decompose aggregate and idiosyncratic volatilities and correlations into high- and low-frequency components using the Factor-Spline-GARCH model proposed by Engle and Rangel (2008) and extended by Rangel and Engle (2009). This semi-parametric approach models the high-frequency component of aggregate and idiosyncratic volatilities as an asymmetric unit GARCH process and the low-frequency component using an exponential quadratic spline. These high- and low-frequency volatilities are combined with idiosyncratic correlations from a dynamic conditional correlation (DCC) model to capture the high- and low-frequency dynamics of conditional correlations. The long-term components reflect movements in volatilities and correlations driven by changes in macroeconomic conditions that affect future cash flows and discount rates. Using a sample of stocks from the Dow Jones Industrial Average (DJIA) index, I document large movements in market volatility, idiosyncratic risk, and market-wide correlations at different horizons. Furthermore, I identify important differences in the cyclical patterns in idiosyncratic volatility across stocks.

These risk dynamics generate significant premia in the cross-section of returns. In particular, I find that long run market volatility risk is priced because of priced innovations in both average long-term idiosyncratic volatilities and average long-term correlations. Both factors carry a significantly negative price of risk. In contrast, short run market volatility risk is only priced in the cross-section because of negatively priced short-term correlation risk, not idiosyncratic risk. The explanatory power of the long and short run variance and correlation factors is not subsumed by traditional size, value, momentum, and liquidity factors, which indicates that they capture different systematic risks. The premia for volatility and correlation risk exhibit strong variation along the value-growth dimension. They are negative for most growth stocks, because these stocks provide insurance against sudden increases in volatilities and correlations. The premia are positive for most value stocks, which perform poorly when aggregate uncertainty increases.

My finding that long- and short-term correlation risk factors are important determinants of time series variation in aggregate returns and cross-sectional variation in individual returns complements recent papers that highlight the role of correlation risk in portfolio choice (Buraschi, Porchia, and Trojani (2009)) and performance evaluation (Buraschi, Kosowski, and Trojani (2009)). Krishnan, Petkova, and Ritchken (2009) also present evidence of a correlation risk premium in stock markets but extract a factor from innovations in historical correlations rather than from revisions in expected future correlations, as required by intertemporal asset pricing models. Furthermore, they do not distinguish between long- and short-term correlation risk and do not examine the pricing of correlation risk at the aggregate level.

The chapter proceeds as follows. Section 4.2 reviews literature on variance and correlation risk. Section 4.3 studies the predictive power of variance and correlation premia for market returns. Section 4.4 provides the theoretical motivation for a cross-sectional premium for variance and correlation risk and describes the methods used to measure long and short run variances and correlations. In Section 4.5 I examine the price of these risk dynamics. Section 4.6 concludes.

4.2 Related Literature on Variance Risk and Correlation Risk

4.2.1 Variance and Correlation Risk Premia in Aggregate Stock Returns

The variance risk premium is formally defined as the difference between the risk-neutral and physical expectation of future variance. Bollerslev, Tauchen, and Zhou (2009) measure the market variance risk premium as the difference between implied and realized market variance and find that it has strong predictive power for aggregate stock market returns at short horizons. This result continues to hold after controlling for traditional return predictors like the dividend yield and the consumption-wealth ratio. They set up a general equilibrium model to explain the predictive power of the variance risk premium. This model extends the long run risk model of Bansal and Yaron (2004) by allowing for time-varying volatility of consumption growth volatility.⁴ The equity premium in their model consists of a consumption risk term as in the traditional consumption CAPM and a volatility risk term. Bollerslev, Tauchen, and Zhou (2009) argue that the variance risk premium isolates the volatility of volatility risk factor. Hence, movements in the variance risk premium reflect changes in variance risk.

An alternative source of variation in the variance risk premium is time-varying risk aversion. The long run risk model and the extension of Bollerslev, Tauchen, and Zhou (2009) assume that risk aversion is constant but Campbell and Cochrane (1999) argue that risk aversion varies over time due to habit formation. Rosenberg and Engle (2002) estimate an empirical pricing kernel and show that their risk aversion measure is countercyclical and positively related to the difference between implied and realized volatility. Bollerslev, Gibson, and Zhou (2009) derive an approximate analytic relation between the variance risk premium and relative risk aversion of the representative investor and link time variation in the premium to macroeconomic variables. Todorov (2009) argues that the market variance risk premium is a compensation for stochastic volatility and price jumps. He documents that even though the effect of jumps on volatility disappears fast, their effect on the variance risk premium is long-lasting, which he attributes to a time-varying attitude of investors towards jumps. Bekaert, Engstrom, and Xing (2009) allow for both time-varying economic uncertainty and time-varying risk aversion. They find that time-varying risk aversion is the dominant source of changes in the equity premium. Bekaert and Engstrom (2009) set up a consumption based model in which preferences are as in the habit formation model of Campbell and Cochrane (1999) but with non-linearities in consumption growth. In their framework the variance premium increases with shocks that generate negative skewness of consumption.

⁴Drechsler and Yaron (2008) extend the long run risk model by allowing for infrequent jumps in long run consumption growth rates and in the volatility of consumption. They argue that changes in the variance premium reflect time variation in agents' perceptions of the risk of these big shocks to the state of the economy, which they want to hedge against.

Market volatility can vary over time due to changes in equity correlations and changes in the volatility of individual stocks. Driessen, Maenhout, and Vilkov (2009) therefore decompose the market variance risk premium into a correlation risk premium and an individual variance premium. Using index options and individual stock options, they find a large variance risk premium for the market index but no evidence of a variance premium for individual stocks. Given the decomposition of stock market volatility, they interpret this as indirect evidence of priced correlation risk. According to this explanation, index options are more expensive than individual options because they can be used to hedge against an increase in market-wide correlations that reduces diversification benefits. An alternative explanation is offered by Buraschi, Trojani, and Vedolin (2009), who link variance and correlation risk premia to market-wide disagreement about future earnings.

If market variance risk is priced because of priced correlation risk, then a testable implication of the theoretical framework of Bollerslev, Tauchen, and Zhou (2009) is that the correlation risk premium should also predict the equity premium. On the other hand, the average individual variance risk premium should not have predictive power for market returns. Hitherto, empirical evidence on the link between idiosyncratic risk and aggregate stock returns is mixed. Goyal and Santa-Clara (2003) take the equal-weighted average stock variance as a measure of idiosyncratic risk and find that it predicts the market return. However, Bali, Cakici, Yan, and Zhang (2005) argue that these results are driven by small NASDAQ stocks and no longer hold when another measure of idiosyncratic risk is used.

4.2.2 Cross-Sectional Evidence on Variance and Correlation Risk

Ang, Hodrick, Xing, and Zhang (2006) find that innovations in aggregate market volatility carry a significantly negative price in the cross-section of stock returns. They argue that investors want to hedge against increases in market volatility because periods of high volatility tend to coincide with a market decline. Thus, stocks with a positive exposure to changes in market volatility provide a hedge against deteriorating investment opportunities. This hedging demand increases the price of those stocks and consequently lowers their expected returns. Using option data, Carr and Wu (2009) find negative variance risk premia for stock indexes and large cross-sectional variation in the variance risk premium for individual stocks. They show that the premium is not explained by traditional risk factors and conclude that it is driven by an independent variance risk factor. Goyal and Saretto (2009) find that a trading strategy that takes a long position in an option portfolio of stocks with a positive difference between realized and implied volatility and a short position in an option portfolio of stocks with a negative difference produces significantly positive alphas. They argue that the difference between implied and realized volatility reflects option mispricing, driven by overreaction to current returns that leads to misestimation of future volatility.⁵

⁵Goyal and Saretto (2009) argue that large deviations between implied and realized volatilities are unlikely to persist because of the high autocorrelation in volatility.

Adrian and Rosenberg (2008) decompose market volatility into long and short run components and find a significantly negative price of risk for both volatility components. They argue that the short run volatility component captures market skewness risk, which they interpret as a measure of the tightness of financial constraints. The long run component is related to business cycle risks. Pricing errors for size and book-to-market sorted portfolios generated by a model that includes the market return and innovations in the two volatility components are smaller than those produced by the Fama and French (1993) three-factor model.

The decomposition of market variance into individual variances and correlations implies that if market variance risk is a priced factor in the cross-section of returns, correlation risk, idiosyncratic risk or both should be priced too. Krishnan, Petkova, and Ritchken (2009) present evidence of a negative correlation risk premium in stock markets, after controlling for asset volatility and other risk factors. However, their correlation factor is based on changes in historical correlations instead of changes in expectations of future correlations. Moreover, they do not distinguish between long- and short-term correlation risk.

Ang, Hodrick, Xing, and Zhang (2006) find a negative cross-sectional relation between idiosyncratic volatility and expected stock returns. However, Fu (2009) finds a positive relation between idiosyncratic volatility and returns when using a conditional measure of idiosyncratic risk. It is important to note that these studies focus on the relation between the *level* of idiosyncratic risk and returns, motivated by the theoretical model of Merton (1987) in which idiosyncratic risk is positively priced in information-segmented markets if investors do not hold well-diversified portfolios. In contrast, an intertemporal asset pricing model predicts a negative price of risk for *innovations* in average individual volatility if an increase in individual volatilities forecasts lower market returns or higher market variances.

4.3 Time Series Analysis of Variance and Correlation Risk

4.3.1 Measurement of Variance and Correlation Risk Premia

Measurement of the variance risk premium requires an estimate of the risk-neutral and physical expectation of future variance. Britten-Jones and Neuberger (2000) demonstrate that the risk-neutral expectation of integrated variance is equal to the model-free implied variance under the assumption that the underlying asset price is continuous but volatility is stochastic. Specifically, they define the model-free implied variance at day s for asset i over the interval Δs as

$$\begin{aligned} IV_{is} &= 2 \int_0^\infty \frac{C_{is}(s + \Delta s, K) - C_{is}(s, K)}{K^2} dK \\ &= 2 \int_0^\infty \frac{C_{is}(s + \Delta s, K) - \max(S_s - K, 0)}{K^2} dK, \end{aligned} \quad (4.3.1)$$

where $C_{is}(T, K)$ is the price at day s of a European call option on asset i with strike price K and maturity at time T and S_s is the price at day s of the underlying. I measure the implied variance for a fixed 30-day maturity, i.e., Δs equals 30, because the options needed to compute the implied variance are most liquid for short maturities. Moreover, the VIX index that is often used by investors as a measure of implied volatility is also defined for a 30-day maturity. I construct the implied variance for each month t , IV_{it} , by taking the average of all daily implied variances within the month, which reduces the noise in the daily measures. Because implied variance is computed directly from option prices, no option pricing model is needed. In contrast, implied variance backed out from the Black and Scholes (1973) model is a flawed estimate of risk-neutral expected variance because the model incorrectly assumes that volatility is constant. Although the model-free implied variance in equation (4.3.1) is defined as the integral over a continuum of strike prices, in practice only a finite number of different strikes is available. However, Jiang and Tian (2005) show that discretization errors are small when the integral is calculated from a limited number of options.

Following Bollerslev, Tauchen, and Zhou (2009), I also use a model-free measure for the realized market variance. In particular, realized variance for asset i over the interval Δt is calculated by summing squared intraday returns

$$RV_{it} = \sum_{j=1}^n \left[p_{it-1+\frac{j}{n}(\Delta)} - p_{it-1+\frac{j-1}{n}(\Delta)} \right]^2, \quad (4.3.2)$$

where p_{it} denotes the logarithmic price of asset i and n is the number of price observations within the interval Δt .⁶

Andersen, Bollerslev, Diebold, and Ebens (2001) show that this approach provides a more precise measure of ex-post latent realized variance than traditional measures of return variation based on returns sampled at a lower frequency. In practice, market microstructure frictions, such as the bid-ask bounce and price discreteness, put an upper limit on the data frequency that can be used to estimate realized variance. While these issues are less important for the market index and for liquid stocks, they can have a serious impact on stocks that are less frequently traded. I therefore follow Driessen, Maenhout, and Vilkov (2009) and restrict the analysis to stocks that are included in the S&P 100 index, because these are actively traded. An additional benefit is that all stocks included in the S&P 100 have exchange-listed options, which are required to compute the implied variance.

The market and individual variance risk premium for a one-month horizon is defined as the difference between the measures of implied variance in (4.3.1) and realized variance in (4.3.2),

$$VRP_{Mt} = IV_{Mt} - RV_{Mt} \quad \text{and} \quad VRP_{it} = IV_{it} - RV_{it}. \quad (4.3.3)$$

⁶Since the realized variance in equation (4.3.2) is measured ex-post, it is strictly speaking not equal to the physical expectation of future variance. However, it has the advantage that it is directly observable, which is important for forecasting purposes. An alternative approach is followed by Bollerslev, Gibson, and Zhou (2009), who rely on a one-factor stochastic volatility model to obtain a forecast of one-step-ahead return variation. Hence, their realized variance measure is no longer model-free and can be affected by model misspecification.

I calculate the average individual variance premium, $\overline{VRP}_t$, as the value-weighted average of the variance premia across all stocks in the market index.

To calculate the monthly correlation risk premium I need an estimate of implied and realized correlations. These measures can be derived by decomposing the variance of the market index into individual variances and correlations,

$$VAR_{Mt} = \sum_{i=1}^{N_t} w_{it}^2 VAR_{it} + \sum_{i=1}^{N_t} \sum_{j \neq i}^{N_t} w_{it} w_{jt} \sqrt{VAR_{it}} \sqrt{VAR_{jt}} COR_{ijt}, \quad (4.3.4)$$

where w_{it} is the index weight of stock i and N_t the number of stocks included in the index at time t .⁷

Because I focus on market-wide correlation risk, I assume that all pairwise correlations are equal, i.e., $COR_{ijt} = COR_t \forall i, j, i \neq j$. Given this assumption and the decomposition of market variance in (4.3.4), average implied and realized correlation at time t can be calculated as⁸

$$IC_t = \frac{IV_{Mt} - \sum_{i=1}^{N_t} w_{it}^2 IV_{it}}{\sum_{i=1}^{N_t} \sum_{j \neq i}^{N_t} w_{it} w_{jt} \sqrt{IV_{it}} \sqrt{IV_{jt}}}, \quad (4.3.5)$$

$$RC_t = \frac{RV_{Mt} - \sum_{i=1}^{N_t} w_{it}^2 RV_{it}}{\sum_{i=1}^{N_t} \sum_{j \neq i}^{N_t} w_{it} w_{jt} \sqrt{RV_{it}} \sqrt{RV_{jt}}}. \quad (4.3.6)$$

The correlation risk premium is defined as the difference between the measures of implied correlation in (4.3.5) and realized correlation in (4.3.6),

$$CRP_t = IC_t - RC_t. \quad (4.3.7)$$

4.3.2 Data Description

Data for S&P 100 index options and for the individual options on the stocks in the S&P 100 index comes from OptionMetrics. The option data is daily and covers the period January 1996 to December 2007. I further retrieve the zero-coupon interest rate curve and underlying stocks prices from CRSP. I apply the same data filters as Driessen, Maenhout, and Vilkov (2009). In particular, I exclude options with zero bid price, zero open interest, missing implied volatility, or missing delta. Furthermore, I remove calls with delta smaller than 0.15 and puts with delta larger than -0.05 because of their high implied volatilities. Since a cross-section of liquid options is needed to calculate model-free implied variances, I focus on short-term

⁷The S&P 100 is a value-weighted index that is rebalanced quarterly. Therefore I compute the weight of each firm based on its market capitalization.

⁸Another way to measure the realized correlation is to compute rolling correlations between all stocks and then calculate the weighted average of all these pairwise correlations. I measure realized correlation using equation (4.3.6) to be consistent with the calculation method for implied correlation. The difference between this estimate of realized correlation and the cross-sectional average of rolling correlations is negligible.

options. Following Carr and Wu (2009), at each day and for each stock I choose the options with the two nearest maturities, except when the shortest maturity is within eight days. In that case I pick the next two maturities to avoid potential microstructure frictions of very short-term options. Moreover, I only consider out-of-the-money and at-the-money options because these options tend to be more liquid than in-the-money options. Finally, I only calculate implied variance on a given day for stocks that have at least two calls and two puts available after the data filtering procedure, to minimize the errors from discretization.

First, I translate option prices into implied volatilities using the Black-Scholes model. I then interpolate these implied volatilities using a cubic spline across moneyness levels, as proposed by Jiang and Tian (2005). To obtain implied volatilities for moneyness levels beyond the available range I extrapolate the implied volatilities at the highest and lowest strike price. This interpolation-extrapolation procedure generates a grid of 1000 implied volatilities for moneyness levels between 0.01 and 3.00. I convert the extracted implied volatilities back into call option prices using the Black-Scholes model and use these call prices to calculate the implied variance in equation (4.3.1). Finally, to obtain the implied variance for a fixed 30-day maturity, I linearly interpolate between the implied variances for the two maturities closest to the 30-day maturity.

To calculate the realized variance in equation (4.3.2), I follow the procedure outlined in Andersen, Bollerslev, Diebold, and Ebens (2001) and use transaction prices from the NYSE trades and quotes (TAQ) database for the components of the S&P 100 and price data from Tickdata for the S&P 100 index. I estimate realized variances using five-minute returns because this sampling frequency provides a reasonable balance between efficiency and robustness to microstructure noise.⁹ For each day I construct equally spaced five-minute returns by computing the logarithmic difference in transaction prices at or immediately before each five-minute mark. To purge the returns from any negative autocorrelation due to the uneven spacing of the observed prices and the bid-ask bounce I de-mean and filter the raw returns using an MA(1) model. Each day I calculate 78 five-minute returns from 9:30am until 16:00pm, including the close-to-open overnight return. Hence, the realized variance estimate for a typical month with 22 trading days is based on $22 \times 78 = 1716$ returns.

Table 4.1 reports summary statistics for the monthly market return, market variance, average individual variance, and average correlation. The mean excess return on the S&P 100 index is 0.42% per month. The average implied market variance is 34.11 per month whereas the average realized market variance is only 20.64, which implies that the variance risk premium equals 13.47 (in percentages squared). In contrast, the value-weighted average implied variance for individual stocks is smaller than the average realized variance, which corresponds to a nega-

⁹The choice of five-minute returns follows Andersen, Bollerslev, Diebold, and Ebens (2001). Driessen, Maenhout, and Vilkov (2009) sample returns at a daily frequency but Bollerslev, Gibson, and Zhou (2009) show that using realized volatilities based on daily returns results in inefficient estimates of the variance risk premium, which loses some of its predictive power for stock returns. I find empirically that using 15-minute returns to calculate realized variances instead of five-minute returns does not affect the main conclusions.

tive variance risk premium on individual stocks. The average implied correlation between all stocks in the index is 0.35 and exceeds the average realized correlation of 0.17, which results in a correlation risk premium of 0.18. The positive premia for market variance risk and correlation risk and the negative premium for individual variance risk are consistent with results documented by Bollerslev, Tauchen, and Zhou (2009) and Driessen, Maenhout, and Vilkov (2009). Interestingly, the correlation risk premium is much more persistent than the market variance premium, although its first-order autocorrelation of 0.84 is lower than that of most traditional return predictors. Panel B shows that the contemporaneous correlation between the market variance premium and the premia for individual variance risk and correlation risk is positive. However, the correlation between the individual variance premium and correlation premium is small and negative, suggesting that these components capture different elements of the aggregate variance premium.

Figure 4.1 plots the implied and realized market variance and the variance premium over the sample period. Both variances and the risk premium increase sharply during periods of market turmoil, such as the Asian crisis in 1997, the LTCM and Russian crisis in 1998, the uncertainty surrounding the war in Iraq around 2003, and the quant crisis in 2007. The average implied and realized variances of individual stocks depicted in Figure 4.2 also rise during these crisis periods. However, the increase in realized variances of individual stocks is much larger than the increase in implied variances, which translates into a negative individual variance risk premium during most crisis periods. Figure 4.3 shows that the correlation risk premium also exhibits strong time variation and increases during turbulent market conditions.

4.3.3 Predicting Stock Market Returns with Variance and Correlation Risk Premia

The finding of Bollerslev, Tauchen, and Zhou (2009) that the market variance premium has predictive power for the equity premium and the conclusion of Driessen, Maenhout, and Vilkov (2009) that market variance risk is only priced because of priced correlation risk, taken together, suggest that the correlation risk premium also predicts market returns. In contrast, the average individual variance risk premium should not have any predictive power.

The first column in Panel A of Table 4.2 confirms that the market variance premium predicts monthly stock market returns, with a t -statistic of 2.33. Column two indicates that the average variance risk premium on individual stocks does not predict market returns. In contrast, the correlation risk premium does have significant predictive power. The adjusted R^2 of the one-month ahead predictability regression including the correlation premium is 3.74%, which is much larger than the explanatory power of the market variance risk premium (1.18%). A one standard deviation increase in the correlation risk premium corresponds to a rise in the predicted market return of more than 10%. Column four shows that the predictive power of the correlation risk premium is not affected by the average individual variance risk premium when both risk premia are jointly added

Table 4.1: Summary Statistics for Variances and Correlations

This table reports descriptive statistics for monthly stock market returns, market variances, individual variances, and correlations for the period from January 1996 to December 2007. R_M is the log return on the S&P 100 in excess of the three-month T-bill rate. IV_M is the model-free implied market variance computed from index options. RV_M is the realized market variance calculated by summing squared five-minute returns on the S&P 100 index over a one-month window. VRP_M is the market variance risk premium, defined as the difference between IV_M and RV_M . $\overline{IV}$ is the value-weighted average model-free implied variance for individual stocks computed from stock options. $\overline{RV}$ is the value-weighted average realized variance of individual stocks calculated by summing squared five-minute stock returns. $\overline{VRP}$ is the variance risk premium on individual stocks, defined as the difference between $\overline{IV}$ and $\overline{RV}$. IC is the implied correlation, calculated from the implied market variance and the implied variance of the stocks in the S&P 100 index. RC is the realized correlation, calculated from the realized market variance and the realized variance of the stocks in the S&P 100 index. CRP is the correlation risk premium, defined as the difference between IC and RC . Panel A reports the mean, median, standard deviation, skewness, kurtosis and autocorrelation for these variables and Panel B presents their pairwise correlations.

	R_M	IV_M	RV_M	VRP_M	$\overline{IV}$	$\overline{RV}$	$\overline{VRP}$	IC	RC	CRP
<i>Panel A: Summary Statistics</i>										
Mean	0.42	34.11	20.64	13.47	95.49	112.64	-17.15	0.35	0.17	0.18
Median	0.91	30.94	15.05	10.48	72.29	94.94	-11.81	0.34	0.16	0.16
Std. Dev.	4.49	24.23	19.15	16.48	61.41	69.43	28.34	0.11	0.08	0.09
Skewness	-0.60	1.88	2.50	1.55	1.09	1.48	-0.77	0.22	1.04	0.22
Kurtosis	3.86	8.73	11.38	13.77	3.29	4.92	4.20	2.58	4.25	2.92
AR(1)	-0.03	0.75	0.65	0.30	0.92	0.77	0.34	0.85	0.70	0.84
<i>Panel B: Correlation Matrix</i>										
R_M	1.00									
IV_M	-0.36	1.00								
RV_M	-0.24	0.74	1.00							
VRP_M	-0.24	0.62	-0.08	1.00						
$\overline{IV}$	-0.19	0.69	0.81	0.07	1.00					
$\overline{RV}$	-0.16	0.64	0.83	-0.03	0.91	1.00				
$\overline{VRP}$	-0.02	-0.06	-0.28	0.24	-0.07	-0.47	1.00			
IC	0.02	0.44	0.20	0.42	-0.08	-0.10	0.07	1.00		
RC	-0.26	0.57	0.67	0.06	0.38	0.24	0.24	0.56	1.00	
CRP	0.25	0.03	-0.34	0.44	-0.42	-0.32	-0.13	0.70	-0.20	1.00

Figure 4.1: Market Variance Risk Premium

This figure plots the implied and realized variance (top panel) and the variance risk premium (bottom panel) for the S&P 100 index from January 1996 to December 2007. The shaded area indicates an NBER recession period.

Figure 4.2: Individual Variance Risk Premium

This figure plots the value-weighted average of the implied and realized variances (top panel) and of the variance risk premia (bottom panel) for the stocks in the S&P 100 index from January 1996 to December 2007. The shaded area indicates an NBER recession period.

Figure 4.3: Correlation Risk Premium

This figure plots the average implied and realized correlation between the stocks in the S&P 100 index (top panel) and the correlation risk premium (bottom panel) from January 1996 to December 2007. The shaded area indicates an NBER recession period.

to the regression. The right side of Panel A shows results for quarterly returns that are based on overlapping monthly observations and scaled by the horizon. I report Hodrick (1992) standard errors that are robust to heteroskedasticity and autocorrelation and account for the overlap in the regressions. The market variance premium and correlation risk premium are also significant in these quarterly regressions. The explanatory power of both regressions is much higher than in the corresponding monthly return regressions. In fact, the adjusted R^2 of the quarterly regression with the correlation risk premium as regressor is almost 15%, which is a direct consequence of its persistence. The individual variance premium remains insignificant.

In Panel B I control for traditional return predictors. Following Bollerslev, Tauchen, and Zhou (2009), I include the default spread, defined as the yield spread between Moody's Baa- and Aaa-rated corporate bonds, the log of the smoothed price-earnings ratio for the S&P 500, defined as the ratio of the price of the index divided by the twelve-month trailing moving average of aggregate earnings, the term spread, defined as the difference between the ten-year and three-month Treasury yield, and the consumption-wealth ratio of Lettau and Ludvigson (2001a).¹⁰ As in Bollerslev, Tauchen, and Zhou (2009), I also include the realized market variance, average individual variance, and average correlation to control for the traditional intertemporal risk-return tradeoff. The coefficient on the market variance premium in the monthly return regressions actually increases after adding the control variables and remains statistically significant. The realized market variance and P/E ratio are also significant and increase the adjusted R^2 by almost 2%. Panel B further shows that the individual variance risk premium is still insignificant when popular predictor variables are added. More importantly, the correlation risk premium stays a significant predictor for one-month ahead returns in the presence of the realized correlation and the traditional predictive variables. Results reported on the right hand side of Panel B show that the market variance and correlation risk premia also remain significant predictors for quarterly returns after controlling for the other predictive variables and the realized risk measures.

Overall, the results from the return predictability regressions show that the predictive power of the market variance risk premium documented by Bollerslev, Tauchen, and Zhou (2009) is driven by the correlation risk premium. Individual variance premia have no predictive power for market returns. This finding supports the claim of Driessen, Maenhout, and Vilkov (2009) that market variance risk is priced because of correlation risk, not individual variance risk.

¹⁰Using other return predictors, like the dividend yield and short-term interest rate, leads to similar results.

Table 4.2: Return Prediction with Variance and Correlation Risk Premia

This table presents results for predictability regressions for the period from January 1996 to December 2007. The dependent variable is the excess return on the S&P 100 index, annualized and in percent. Results on the left are for a one-month horizon and estimates on the right correspond to quarterly returns that are scaled by the horizon. The quarterly regressions are based on overlapping monthly returns. In Panel A the independent variables are VRP_M , $\overline{VRP}$, and CRP , all defined in Table 4.1. Panel B adds traditional return predictors and realized risk measures as independent variables to the regressions. RV_M , $\overline{RV}$, and RC are defined in Table 4.1. CAY is the consumption-wealth ratio, for which the most recently available quarterly observations are taken as monthly observations, DEF is the default spread, defined as the yield differential between bonds rated Baa by Moody's and bonds with a Moody's rating of Aaa, $\log(P/E)$ is the log of the smoothed price-earnings ratio for the S&P 500, and $TERM$ is the term spread, defined as the difference between the ten-year and three-month Treasury yield. DEF and $\log(P/E)$ are multiplied by 12. Intercept estimates are not reported to save space. t -Statistics based on Hodrick (1992) standard errors are in parentheses.

	Monthly Return Regressions				Quarterly Return Regressions			
<i>Panel A: Predictability Regressions Excluding Traditional Predictors</i>								
VRP_{Mt-1}	0.45 (2.33)				0.52 (4.66)			
$\overline{VRP}_{t-1}$		-0.09 (-0.53)		-0.04 (-0.28)		-0.15 (-1.18)		-0.10 (-1.03)
CRP_{t-1}			120.04 (2.81)	118.36 (2.73)			125.34 (3.88)	121.41 (3.86)
Adj. R^2 (%)	1.18	-0.46	3.74	3.11	7.46	0.01	14.88	15.19
<i>Panel B: Predictability Regressions Including Traditional Predictors</i>								
VRP_{Mt-1}	0.60 (2.93)				0.67 (4.16)			
$\overline{VRP}_{t-1}$		-0.04 (-0.16)		-0.03 (-0.10)		-0.19 (-1.02)		-0.19 (-1.26)
CRP_{t-1}			140.26 (2.01)	147.26 (2.06)			156.97 (2.84)	146.73 (3.07)
RV_{Mt-1}	0.43 (1.98)				0.33 (2.21)			
$\overline{RV}_{t-1}$		-0.04 (-0.34)		-0.02 (-0.13)		-0.10 (-1.24)		-0.08 (-0.78)
RC_{t-1}			94.43 (1.98)	113.31 (2.02)			126.83 (2.44)	132.84 (2.26)
CAY_{t-1}	2.52 (1.04)	6.06 (2.53)	1.19 (0.43)	1.56 (0.37)	0.60 (0.33)	4.36 (2.57)	-1.20 (-0.55)	-0.12 (-0.04)
DEF_{t-1}	-0.28 (-0.16)	0.40 (0.21)	-0.25 (-0.13)	-0.30 (-0.13)	-1.42 (-0.96)	-0.63 (-0.39)	-1.28 (-0.78)	-1.33 (-0.71)
$\log(P/E)_{t-1}$	-4.65 (-2.19)	-1.53 (-0.58)	-1.55 (-0.91)	-1.23 (-0.35)	-3.97 (-2.33)	0.04 (0.02)	-0.82 (-0.55)	0.35 (0.15)
$TERM_{t-1}$	1.56 (0.33)	-2.18 (-0.48)	-6.75 (-1.59)	-7.21 (-1.42)	2.83 (0.71)	-1.62 (-0.45)	-5.42 (-1.50)	-6.62 (-1.72)
Adj. R^2 (%)	2.97	0.02	4.68	3.48	15.75	6.56	21.95	22.07

4.4 Long and Short Run Variances and Correlations

4.4.1 ICAPM with Time-Varying Variances and Correlations

The existence of a market variance risk premium and correlation risk premium in index options, documented in the previous section, implies that investors are willing to pay a premium to insure against shocks to systematic volatility and market-wide correlations. In this section I study the pricing implications of investor aversion to variance and correlation risk for the cross-section of stock returns. The theoretical motivation for a cross-sectional risk premium for variance and correlation risk is given by the intertemporal capital asset pricing model (ICAPM) of Merton (1973). The main premise of this model is that investors want to hedge against a deterioration of investment opportunities. Intuitively, assets that perform poorly when the market return is low, when expected future market returns fall, or when expected future market variances increase should have higher expected returns. Consequently, in this model any variable that forecasts future market returns or variances is a relevant state variable of the pricing kernel.

Because future market variance can change due to variation in future individual variances and future correlations, the ICAPM implies that shocks to average individual variances, innovations in market-wide correlations, or both should be priced. I test this prediction by estimating the following regression equation,

$$r_{it+1} = \alpha_i + \gamma Cov_t(r_{it+1}, r_{Mt+1}) + \lambda' Cov_t(r_{it+1}, S_{t+1}) + \epsilon_{it+1}, \quad (4.4.1)$$

where γ is the coefficient of relative risk aversion and where the elements in λ are the prices of the hedge-related risks in the vector S_{t+1} . The conditional covariances in (4.4.1) measure the sensitivity of asset i to movements in the market return and in the innovations in the state variables. Since my objective is to study the pricing of variance and correlation risk, I include in S_{t+1} shocks to conditional market volatility, average idiosyncratic volatility, and market-wide correlations.¹¹

Although these volatilities and correlations are exogenous in the intertemporal model in (4.4.1), it is important to identify their underlying economic determinants to shed more light on the sources of variance and correlation risk premia. In an efficient market, movements in asset prices reflect revisions in the expected value of discounted dividends. These changes in expectations are driven by new information about the prospects of the firm and the expected return used to discount future cash flows. Because both the intensity and impact of this news varies over time and across firms, volatilities and correlations will fluctuate. Campbell (1991) formalizes this idea by decomposing unexpected returns into revisions in expected dividend growth (cash flow news) and revisions in expected returns (discount rate news). Following this approach, the conditional variance of returns can

¹¹Goyal and Santa-Clara (2003) show that idiosyncratic risk constitutes a very large fraction of the average individual stock variance.

be decomposed into

$$Var_t(r_{it+1}) = Var_t(N_{it+1}^{CF}) + Var_t(N_{it+1}^{DR}) - 2Cov_t(N_{it+1}^{CF}, N_{it+1}^{DR}). \quad (4.4.2)$$

Thus, return variance depends on the variance of cash flow news and discount rate news. Using the unconditional version of (4.4.2), Campbell (1991) finds that most stock market volatility is due to discount rate news. In contrast, Vuolteenaho (2002) shows that most variation in individual stock returns is due to cash flow news. These findings can be reconciled by noticing that most cash flow news at the firm level is idiosyncratic and cancels out at the market level.

Similarly, Engle (2009a) shows that the conditional covariance can be written as

$$\begin{aligned} Cov_t(r_{it+1}, r_{jt+1}) = & Cov_t(N_{it+1}^{CF}, N_{jt+1}^{CF}) + Cov_t(N_{it+1}^{DR}, N_{jt+1}^{DR}) \\ & - Cov_t(N_{it+1}^{CF}, N_{jt+1}^{DR}) - Cov_t(N_{it+1}^{DR}, N_{jt+1}^{CF}). \end{aligned} \quad (4.4.3)$$

Consequently, in this framework return correlations depend on the correlations between dividend news of assets and between their discount rate news. Because cash flow news and discount rate news can have a different impact in the long run than in the short run, asset volatilities and correlations will exhibit high- and low-frequency movements. Consistent with this intuition, several studies find empirically that two component volatility specifications better explain equity volatility than single-factor models (see, for instance, Engle and Lee (1999) and Chernov, Gallant, Ghysels, and Tauchen (2003)). As a result, shocks to long run variances and correlations can be priced differently from shocks to short run variances and correlations and should be included as separate factors in the intertemporal model in (4.4.1).

4.4.2 Specification of Volatility and Correlation Components

To study these cross-sectional predictions, I separate systematic and idiosyncratic volatilities and correlations into long- and short-term components using the Factor-Spline-GARCH-DCC (FSG-DCC) model first proposed by Engle and Rangel (2008) for volatilities and subsequently extended by Rangel and Engle (2009) for correlations.¹² This semi-parametric approach models the high-frequency component of aggregate and idiosyncratic volatilities as an asymmetric GARCH process and the low-frequency component using an exponential quadratic spline. These high- and low-frequency volatilities are then combined with idiosyncratic

¹²An alternative approach to split correlations into long- and short-term components is suggested by Colacito, Engle, and Ghysels (2009). They specify a DCC-MIDAS model, in which short run fluctuations in correlations are captured by a standard DCC scheme and where the long run component is a weighted average of historical correlations. The advantage of this approach is that it estimates the long and short run correlations directly. However, because the weights need to be estimated, this approach is more prone to misspecification than the nonparametric approach to approximate long-term volatilities and correlations taken by Rangel and Engle (2009).

correlations from a standard dynamic conditional correlation (DCC) model to capture the high- and low-frequency dynamics of conditional correlations. The FSG-DCC model is a reduced form model purely designed to measure long and short run volatilities and correlations and as such it is not directly connected to stylized general equilibrium models such as those proposed by Bollerslev, Tauchen, and Zhou (2009) and Drechsler and Yaron (2008). Nevertheless, Engle and Rangel (2008) and Rangel and Engle (2009) do link the long-term volatility components to macroeconomic fundamentals that affect future cash flows and discount rates, such as short-term interest rates, GDP growth, and inflation.

Following Engle and Rangel (2008), I capture long and short run patterns in systematic volatility by writing market returns as

$$r_{Mt} = \alpha_M + u_{Mt} = \alpha_M + \sqrt{\tau_{Mt}g_{Mt}}\epsilon_{Mt} \quad \text{with} \quad \epsilon_{Mt}|\Phi_{t-1} \sim (0, 1), \quad (4.4.4)$$

where r_{Mt} is the excess market return, α_M the unconditional market return, u_{Mt} the market return innovation, and Φ_{t-1} the information set up to time $t - 1$. g_{Mt} and τ_{Mt} are the high- and low-frequency components of market volatility, respectively. This multiplicative decomposition of conditional volatility allows both the intensity of news and its impact on volatility to vary over time, for instance depending on the state of the economy. g_{Mt} reflects transitory volatility effects that could be related to liquidity shocks and τ_{Mt} captures long-term movements in volatility driven by macroeconomic conditions that affect future cash flows and discount rates.¹³

The high-frequency component g_{Mt} is modeled as a unit GARCH process,

$$\begin{aligned} g_{Mt} = & \left(1 - \theta_M - \phi_M - \frac{\gamma_M}{2}\right) + \theta_M \frac{(r_{Mt-1} - \alpha_M)^2}{\tau_{Mt-1}} \\ & + \gamma_M \frac{(r_{Mt-1} - \alpha_M)^2}{\tau_{Mt-1}} I_{r_{Mt-1} < 0} + \phi_M g_{Mt-1}, \end{aligned} \quad (4.4.5)$$

where θ_M is the ARCH effect, ϕ_M the GARCH coefficient, and $I_{r_{Mt-1} < 0}$ an indicator function for negative market returns. The γ_M coefficient in this mean-reverting specification captures the leverage effect observed in stock market data where negative returns have a stronger impact on volatility than positive returns. I ensure stationarity in the GARCH process by imposing the restrictions that θ_M , ϕ_M , γ_M , and $(1 - \theta_M - \phi_M - \frac{\gamma_M}{2})$ should be positive.

The low-frequency component τ_{Mt} is modeled as an exponential spline,

$$\tau_{Mt} = c_M \exp \left(w_{M0}t + \sum_{k=1}^{K_M} w_{Mk}((t - t_{Mk-1})_+)^2 \right), \quad (4.4.6)$$

where $(t - t_{Mk-1})_+ = (t - t_{Mk-1})$ if $t > t_{Mk-1}$ and 0 otherwise. c_M is the low-frequency volatility at $t = 0$. K_M denotes the number of equally spaced knots

¹³Engle, Ghysels, and Sohn (2009) also specify a two-component volatility model but directly link the long-term volatility component to macroeconomic state variables using a mixed data sampling (MIDAS) approach.

and determines the number of cycles in low-frequency volatility. Following Engle and Rangel (2008), I choose the optimal number of knots based on the Bayesian Information Criterion (BIC). The w_{Mk} coefficients determine the curvature of the volatility cycles. Because g_{Mt} is normalized to have an unconditional mean equal to one, unconditional volatility is equal to τ_{Mt} in this model. The Spline-GARCH model therefore allows unconditional volatility to slowly vary over time, unlike traditional GARCH implementations in which it is constant. The exponential form of the spline ensures that low-frequency volatility is always positive.

I specify a single-factor model to capture the covariance matrix of individual stock returns in a parsimonious way and to split returns into systematic and idiosyncratic components,

$$r_{it} = \alpha_i + \beta_i r_{Mt} + u_{it}, \quad (4.4.7)$$

where r_{it} is the excess return on stock i , α_i the pricing error, and β_i the market beta. The one-factor structure imposes a restriction on the covariance matrix by assuming that the idiosyncratic return u_{it} is uncorrelated with the market return and uncorrelated across stocks. If these moment conditions hold conditionally, the conditional correlation between two assets can be written as

$$\begin{aligned} Cor_{t-1}(r_{it}, r_{jt}) = & \frac{\beta_i \beta_j Var_{t-1}(r_{Mt})}{\sqrt{\beta_i^2 Var_{t-1}(r_{Mt}) + Var_{t-1}(u_{it})} \sqrt{\beta_j^2 Var_{t-1}(r_{Mt}) + Var_{t-1}(u_{jt})}}. \end{aligned} \quad (4.4.8)$$

Since betas are assumed to be constant in this expression, any time variation in conditional correlations is completely driven by variation in conditional market variance and idiosyncratic variances. However, a large body of empirical evidence shows that betas vary over time as a function of firm-specific and macroeconomic conditions. When true betas are time-varying, the assumption that the idiosyncratic return is conditionally uncorrelated with the market return no longer holds. Furthermore, it is likely that the single-factor structure in equation (4.4.7) is too simplistic and that latent factors affect asset returns and correlations. In that case, the restriction that idiosyncratic returns are cross-sectionally uncorrelated is violated.

To capture these important empirical features of the data and allow for richer correlation dynamics, Engle (2009b) allows for temporal deviations from these restrictions. In this extension, unobserved factors can influence conditional correlations and betas can vary over time but should be covariance-stationary and mean-revert to their unconditional expectation.¹⁴ This generalization yields the

¹⁴In particular, while the assumptions that the idiosyncratic return is uncorrelated with the market return and uncorrelated across stocks should still hold unconditionally, in the generalization of Engle (2009b) these restrictions no longer need to be satisfied conditionally. See Rangel and Engle (2009) for a detailed proof.

following expression for the conditional correlation

$$\begin{aligned}
 Cor_{t-1}(r_{it}, r_{jt}) = & \left[\beta_i \beta_j Var_{t-1}(r_{Mt}) + \beta_i E_{t-1}(r_{Mt} u_{jt}) \right. \\
 & \left. + \beta_j E_{t-1}(r_{Mt} u_{it}) + E_{t-1}(u_{it} u_{jt}) \right] \\
 & \times \left[\sqrt{\beta_i^2 Var_{t-1}(r_{Mt}) + Var_{t-1}(u_{it}) + 2\beta_i E_{t-1}(r_{Mt} u_{it})} \right. \\
 & \left. \times \sqrt{\beta_j^2 Var_{t-1}(r_{Mt}) + Var_{t-1}(u_{jt}) + 2\beta_j E_{t-1}(r_{Mt} u_{jt})} \right]^{-1}.
 \end{aligned} \tag{4.4.9}$$

The second and third term in the numerator and the terms added to the denominator capture the effects of time-varying betas on conditional correlations. The impact of latent factors is reflected in the last term of the numerator. The empirical implementation of (4.4.9) and decomposition of conditional correlations into high- and low-frequency terms requires estimates of long and short run market volatility and idiosyncratic volatilities and estimates of correlations between the market return and idiosyncratic returns and among the idiosyncratic returns themselves.

The specification of high- and low-frequency components of idiosyncratic volatility is similar to the decomposition of market volatility. In particular, I model the idiosyncratic return u_{it} as

$$\begin{aligned}
 u_{it} &= \sqrt{\tau_{it} g_{it}} \epsilon_{it} \quad \text{with} \quad \epsilon_{it} | \Phi_{t-1} \sim (0, 1) \\
 g_{it} &= \left(1 - \theta_i - \phi_i - \frac{\gamma_i}{2} \right) + \theta_i \frac{(r_{it-1} - \alpha_i - \beta_i r_{Mt-1})^2}{\tau_{it-1}} \\
 &\quad + \gamma_i \frac{(r_{it-1} - \alpha_i - \beta_i r_{Mt-1})^2}{\tau_{it-1}} I_{r_{it-1} < 0} + \phi_i g_{it-1} \\
 \tau_{it} &= c_i \exp \left(w_{i0} t + \sum_{k=1}^{K_i} w_{ik} ((t - t_{ik-1})_+)^2 \right) \quad \forall i,
 \end{aligned} \tag{4.4.10}$$

where g_{it} and τ_{it} are the high- and low-frequency components of idiosyncratic volatility. Because the intensity of news and its effect on idiosyncratic risk depends on firm-specific conditions, I allow the number of cycles in long-term idiosyncratic volatility to vary across stocks based on the BIC.

I estimate the conditional correlations among the idiosyncratic returns and between these shocks and the market return using the DCC approach of Engle (2002).¹⁵ The standardized return innovations that are used as inputs in the DCC model are computed using the conditional volatilities from the univariate

¹⁵Engle (2009a) points out that when the factor model in equation (4.4.7) is correctly specified (i.e., beta is constant and there are no latent factors), all correlations produced by the DCC model should be zero.

Spline-GARCH models,

$$\epsilon_{Mt} = \frac{u_{Mt}}{\sqrt{\tau_{Mt}g_{Mt}}} \quad \text{and} \quad \epsilon_{it} = \frac{u_{it}}{\sqrt{\tau_{it}g_{it}}}. \quad (4.4.11)$$

The standard mean-reverting DCC specification is given by

$$\begin{aligned} \rho_{i,j,t}^\epsilon &= \frac{q_{i,j,t}}{\sqrt{q_{i,i,t}}\sqrt{q_{j,j,t}}}, \\ q_{i,j,t} &= \bar{\rho}_{i,j}^\epsilon + \alpha_{DCC}(\epsilon_{it-1}\epsilon_{jt-1} - \bar{\rho}_{i,j}^\epsilon) + \beta_{DCC}(q_{i,j,t-1} - \bar{\rho}_{i,j}^\epsilon) \quad \forall i, j, \\ q_{M,i,t} &= \bar{\rho}_{M,i}^\epsilon + \alpha_{DCC}(\epsilon_{Mt-1}\epsilon_{it-1} - \bar{\rho}_{M,i}^\epsilon) + \beta_{DCC}(q_{M,i,t-1} - \bar{\rho}_{M,i}^\epsilon) \quad \forall i, \end{aligned} \quad (4.4.12)$$

where $q_{i,j,t}$ and $q_{M,i,t}$ are quasi-correlations, $\bar{\rho}_{i,j}^\epsilon = E(\epsilon_{it}\epsilon_{jt})$ because the return innovations are standardized to have variance equal to one, and $\bar{\rho}_{i,i}^\epsilon = 1$. I take the sample correlations as estimators of the intercept parameters $\bar{\rho}_{i,j}^\epsilon$. Furthermore, given the assumptions on the nature of time-varying betas (footnote 14), $\bar{\rho}_{M,i}^\epsilon = 0 \quad \forall i$. Positive definiteness of the quasi-correlation matrix is guaranteed by imposing the constraints that α_{DCC} , β_{DCC} , and $(1 - \alpha_{DCC} - \beta_{DCC})$ are positive.

Substituting the conditional expectations from the specifications of high- and low-frequency market volatility and idiosyncratic volatility and the idiosyncratic correlations from the DCC model into the expression for conditional correlations in equation (4.4.9), Rangel and Engle (2009) show that the high-frequency conditional correlation can be written as

$$\begin{aligned} Cor_{t-1}(r_{it}, r_{jt}) &= \left[\beta_i \beta_j \tau_{Mt} g_{Mt} + \beta_i \sqrt{\tau_{Mt} g_{Mt}} \sqrt{\tau_{jt} g_{jt}} \rho_{M,j,t}^\epsilon \right. \\ &\quad \left. + \beta_j \sqrt{\tau_{Mt} g_{Mt}} \sqrt{\tau_{it} g_{it}} \rho_{M,i,t}^\epsilon + \sqrt{\tau_{it} g_{it}} \sqrt{\tau_{jt} g_{jt}} \rho_{i,j,t}^\epsilon \right] \\ &\quad \times \left[\sqrt{\beta_i^2 \tau_{Mt} g_{Mt} + \tau_{it} g_{it} + 2\beta_i \sqrt{\tau_{Mt} g_{Mt}} \sqrt{\tau_{it} g_{it}} \rho_{M,i,t}^\epsilon} \right. \\ &\quad \left. \times \sqrt{\beta_j^2 \tau_{Mt} g_{Mt} + \tau_{jt} g_{jt} + 2\beta_j \sqrt{\tau_{Mt} g_{Mt}} \sqrt{\tau_{jt} g_{jt}} \rho_{M,j,t}^\epsilon} \right]^{-1}. \end{aligned} \quad (4.4.13)$$

Substitution of unconditional expectations from the expressions for volatilities and correlations into the unconditional version of (4.4.9) gives the low-frequency component of this correlation

$$\overline{Cor}_{t-1}(r_{it}, r_{jt}) = \frac{\beta_i \beta_j \tau_{Mt} + \sqrt{\tau_{it}} \sqrt{\tau_{jt}} \bar{\rho}_{i,j}^\epsilon}{\sqrt{\beta_i^2 \tau_{Mt} + \tau_{it}} \sqrt{\beta_j^2 \tau_{Mt} + \tau_{jt}}}. \quad (4.4.14)$$

Equation (4.4.14) shows that fluctuations in long run correlations reflect trends in long-term systematic and idiosyncratic variances. The high-frequency correlation in equation (4.4.13) mean-reverts to this time-varying low-frequency component. In contrast, in the standard mean-reverting DCC model conditional correlations mean-revert to constant unconditional correlations.

4.4.3 Estimation of Long and Short Run Volatilities and Correlations

I use a two-step approach to estimate the FSG-DCC model. The first stage estimates the alphas and betas of the factor model and the parameters of the univariate Spline-GARCH models for each stock and for the market index. I assume a Student's t -distribution for the return innovations because the distribution of stock returns typically exhibits fat tails.¹⁶ The log-likelihood function for the Spline-GARCH models with t -distributed innovations is given by

$$\begin{aligned} L(\eta_M, \nu_M) &= \log \left(\frac{\Gamma((\nu_M + 1)/2)}{\Gamma(\nu_M/2) \sqrt{(\nu_M - 2) \tau_{Mt} g_{Mt}}} \right) \\ &\quad - \frac{(\nu_M + 1)}{2} \log \left(1 + \frac{(r_{Mt} - \alpha_M)^2}{\tau_{Mt} g_{Mt} (\nu_M - 2)} \right) \\ L(\eta_i, \nu_i) &= \log \left(\frac{\Gamma((\nu_i + 1)/2)}{\Gamma(\nu_i/2) \sqrt{(\nu_i - 2) \tau_{it} g_{it}}} \right) \\ &\quad - \frac{(\nu_i + 1)}{2} \log \left(1 + \frac{(r_{it} - \alpha_i - \beta_i r_{Mt})^2}{\tau_{it} g_{it} (\nu_i - 2)} \right), \end{aligned} \quad (4.4.15)$$

where η is a vector that contains the alphas, betas, and volatility parameters of the Spline-GARCH models. Γ denotes the gamma function with ν degrees of freedom.

The second step estimates the DCC parameters using the standardized return innovations from the first stage as input. The log-likelihood for the DCC model is

$$L(\eta, \psi) = -\frac{1}{2} \sum_{t=1}^T (\log |R_t| + \epsilon'_t R_t^{-1} \epsilon_t), \quad (4.4.16)$$

where ψ is a vector that contains the DCC parameters and R_t is the conditional correlation matrix of the idiosyncratic returns and the market return. This second log-likelihood function is maximized conditional on the mean and variance estimates from the first stage.

I estimate the high- and low-frequency variances and correlations for the S&P 500 index and for the 30 components of the Dow Jones Industrial Average (DJIA) index as of February 18, 2008.¹⁷ The sample is daily and covers the period from

¹⁶Alternatively, a Gaussian quasi-maximum likelihood approach can be used. However, Rangel and Engle (2009) point out that assuming normality might lead to inaccurate choices of the number of knots in the spline function and, consequently, to misspecified volatilities. I find empirically that for most stocks the Gaussian QML approach leads to the selection of a larger number of knots than the Student's t -distributional assumption. Nevertheless, our main empirical results are not affected by this choice.

¹⁷I restrict the analysis to the 30 Dow stocks because even though the FSG-DCC model is very parsimonious, the estimation takes a long time for larger covariance matrices since the maximum likelihood estimator needs to invert a $N \times N$ matrix many times. Furthermore, the stocks in the DJIA are liquid and represent an important part of the aggregate stock market. I use the S&P 500 index as a proxy for the market because of its broader market coverage than the S&P 100. However, using the S&P 100 index leads to similar results.

January 1990 to December 2008. Table 4.3 presents the average excess return and unconditional standard deviation for these stocks. The IT stocks Microsoft and Intel have the highest return while General Motors (GM) and American International Group (AIG) have the lowest returns. AIG and GM are also the most volatile stocks, largely driven by their financial distress during the credit crisis in 2007 and 2008.

Table 4.3: Descriptive Statistics for Dow 30 Components

This table reports the average daily excess return and unconditional standard deviation for the stocks in the Dow 30 index after the index revision in February 2008. The sample covers the period from January 1990 to December 2008.

Ticker	Company Name	Mean Return (%)	Std. Dev. (%)
AA	Alcoa	0.023	2.31
AIG	American International Group	-0.007	2.77
AXP	American Express	0.032	2.25
BA	Boeing	0.027	1.97
BAC	Bank of America	0.031	2.32
C	Citigroup	0.051	2.65
CAT	Caterpillar	0.051	2.03
CVX	Chevron	0.042	1.60
DD	DuPont	0.017	1.79
DIS	Walt Disney	0.027	2.01
GE	General Electric	0.033	1.75
GM	General Motors	-0.015	2.66
HD	Home Depot	0.065	2.20
HPQ	Hewlett-Packard	0.064	2.55
IBM	International Business Machines	0.036	1.93
INTC	Intel	0.078	2.69
JNJ	Johnson & Johnson	0.047	1.51
JPM	JP Morgan	0.053	2.46
KO	Coca-Cola	0.036	1.57
MCD	McDonald's	0.045	1.71
MMM	3M	0.029	1.50
MRK	Merck	0.031	1.84
MSFT	Microsoft	0.085	2.21
PFE	Pfizer	0.049	1.85
PG	Procter & Gamble	0.046	1.57
T	AT&T	0.028	1.81
UTX	United Technologies	0.052	1.78
VZ	Verizon	0.021	1.74
WMT	Wal Mart	0.053	1.86
XOM	Exxon Mobil	0.047	1.56

Panel A in Table 4.4 reports parameter estimates for the FSG-DCC model for the individual stocks and the S&P 500 index.¹⁸ Estimates of the ARCH effects are presented in column 3. All ARCH coefficients are significant at the 5% level except those of Intel and Coca-Cola, which are significant at the 10% level. Individual ARCH effects range from 0.01 to 0.13 with a mean of 0.06. In contrast, the ARCH effect for the market factor is close to zero and insignificant. Column 4 shows the GARCH estimates, which are significant for all stocks and the market factor at a 1% level. The GARCH effect is between 0.22 and 0.98 for individual stocks with a mean value of 0.77 and equals 0.90 for the market index. As pointed out by Rangel and Engle (2009), the large cross-sectional variation in ARCH and GARCH effects indicates that there is substantial variation in the persistence of idiosyncratic volatilities across stocks.

Estimates of the leverage effects are in column 5. The leverage coefficient is significant for approximately half of the individual stocks and ranges from 0.00 to 0.12 with a mean of 0.05. The leverage effect for the market index is much stronger than for most individual stocks. Column 6 reports the estimated degrees of freedom of the Student's t -distribution, which are between 4 and 9 with a mean of 6. These values indicate excess kurtosis, which highlights the importance of assuming a t -distribution for return innovations in the estimation instead of the normal distribution. The last column shows the optimal number of knots in the spline function based on the BIC. This number varies from 1 to 9 for the individual stocks and is equal to 6 for the market. Since the sample covers 19 years and the knots are equally spaced, this means that the length of each market volatility cycle is just over 3 years. The large variation in the number of knots across firms reveals important differences in the cyclical patterns in their idiosyncratic volatilities.

Panel B in Table 4.4 presents estimates of the second-stage DCC parameters, which are both significant at a 1% level. α_{DCC} equals 0.0028 and β_{DCC} is 0.9920, indicating strong persistence in the correlation between idiosyncratic returns.

To verify whether the FSG-DCC model produces reasonable correlation estimates, in Figure 4.4 I plot the equal-weighted cross-sectional average of the conditional correlations from the FSG-DCC model and of the 100-day historical correlations between the 30 Dow Jones stocks. The correlations are converted to a monthly frequency by taking the average of all daily correlations in each month. The patterns in both correlation measures are very similar, which confirms that the FSG-DCC model adequately captures the dynamics of average correlations. The plot further shows that movements in historical correlations lag changes in conditional correlations, consistent with the notion that the latter are expectations of one-month ahead correlations.

Figures 4.5 and 4.6 show the annualized high- and low-frequency components of market volatility and average idiosyncratic volatility, respectively. Both graphs reveal distinct cyclical patterns in volatility and highlight the importance of allowing unconditional (low-frequency) volatility to vary over time. Consistent with

¹⁸I do not report alphas and betas to save space. The alphas are generally small and only significantly different from zero at a 5% level for two stocks, AIG and GM. As expected, market betas are centered around one and range from 0.62 for Johnson & Johnson to 1.39 for Intel.

Table 4.4: **Spline-GARCH-DCC Estimates**

This table reports estimation results of the Spline-GARCH model in Panel A and the DCC model in Panel B for the 30 stocks in the DJIA and for the S&P 500 index. For stock i , θ_i is the ARCH effect, ϕ_i the GARCH effect, γ_i the leverage effect, ν_i the degrees of freedom of the Student's t -distribution, and K_i the optimal number of knots in the spline function. α_{DCC} and β_{DCC} are the DCC parameters. The sample is daily and covers the period January 1990 to December 2008.

	(1)		(2)		(3)		(4)	(5)
	θ_i	t -stat	ϕ_i	t -stat	γ_i	t -stat	ν_i	K_i
<i>Panel A: Spline-GARCH Estimates</i>								
AA	0.02	3.47	0.94	62.17	0.02	2.03	6	4
AIG	0.06	4.08	0.82	30.26	0.07	3.61	7	7
AXP	0.05	3.66	0.83	23.87	0.06	3.25	6	6
BA	0.04	2.36	0.84	10.78	0.06	2.54	6	4
BAC	0.12	5.40	0.65	13.08	0.08	2.66	6	8
C	0.09	4.31	0.74	15.98	0.09	3.46	6	4
CAT	0.08	3.98	0.40	2.94	0.00	0.04	5	9
CVX	0.05	4.90	0.91	55.57	0.02	1.16	9	3
DD	0.07	2.94	0.75	6.65	0.00	0.03	6	5
DIS	0.10	3.89	0.56	6.25	0.04	1.25	5	4
GE	0.04	2.61	0.84	15.30	0.06	3.05	7	4
GM	0.02	2.53	0.96	107.59	0.03	4.52	5	4
HD	0.02	2.74	0.91	32.71	0.05	3.64	6	4
HPQ	0.01	3.85	0.98	214.39	0.00	0.24	4	4
IBM	0.07	3.32	0.68	11.38	0.12	3.52	4	6
INTC	0.02	1.90	0.93	42.28	0.03	2.61	5	4
JNJ	0.04	3.56	0.83	29.27	0.12	4.85	6	4
JPM	0.03	2.51	0.88	23.41	0.07	3.84	6	5
KO	0.04	1.75	0.81	17.51	0.10	3.35	6	4
MCD	0.03	4.84	0.96	165.41	0.01	0.64	6	1
MMM	0.13	4.57	0.22	2.44	0.00	0.01	5	8
MRK	0.09	3.26	0.44	4.56	0.09	2.18	5	8
MSFT	0.08	4.36	0.77	19.45	0.06	2.19	5	4
PFE	0.08	5.63	0.81	22.41	0.01	0.91	6	5
PG	0.05	3.09	0.79	14.18	0.06	2.52	6	8
T	0.07	4.71	0.84	26.51	0.01	0.87	7	4
UTX	0.08	3.55	0.60	5.84	0.08	2.35	6	6
VZ	0.11	4.50	0.67	7.72	0.01	0.83	7	4
WMT	0.03	5.40	0.96	150.16	0.00	0.52	6	1
XOM	0.08	5.24	0.86	35.36	0.00	0.18	8	3
SP500	0.00	0.61	0.90	70.90	0.13	8.31	9	6
	α_{DCC}	t -stat	β_{DCC}	t -stat				
<i>Panel B: DCC Estimates</i>								
DCC	0.0028	16.51	0.99	1331.41				

Figure 4.4: **Average Conditional and Historical Correlation**

This figure plots the equal-weighted average of conditional correlations and of 100-day rolling correlations between 30 Dow Jones stocks for the period January 1990 to December 2008. The conditional correlations are obtained from the Factor-Spline-GARCH-DCC model outlined in Section 4.4.2 and the estimation is based on daily returns. The correlations are converted to a monthly frequency by taking the average of all daily correlations in each month. Shaded areas indicate NBER recession periods.

Campbell, Lettau, Malkiel, and Xu (2001), I observe an upward trend in idiosyncratic risk until 2001. After 2001, however, idiosyncratic volatility decreases until 2006 and starts to rise again from 2007 onwards due to the credit crisis, which is in line with results of Bekaert, Hodrick, and Zhang (2009b). In contrast to the increase in idiosyncratic volatility at the beginning of the sample, low-frequency market volatility decreases until 1997. After 1997, movements in market volatility more closely resemble those in idiosyncratic volatility.

The trends in market volatility and average idiosyncratic volatility suggest that correlations also exhibit long-term trends. In particular, the one-factor structure in (4.4.7) implies that when market and idiosyncratic risk move in the same direction, they have opposite effects on correlations. An increase in market volatility leads to an increase in correlations whereas an increase in idiosyncratic risk results in a decrease in correlations. In contrast, when market and idiosyncratic volatility change in opposite directions, correlations can exhibit large movements. These patterns in market-wide correlations are illustrated in Figure 4.7. The smooth nature of the long-term volatilities and correlations is consistent with the dynamics of the slow-moving fundamental economic variables that are associated with these

Figure 4.5: Long and Short Run Market Volatility

This figure plots the high- and low-frequency components of market volatility for the period January 1990 to December 2008. The high-frequency component is modeled as a unit GARCH process and the low-frequency component is modeled as an exponential quadratic spline. The estimation is based on daily returns on the S&P 500 index and the resulting volatilities are annualized.

components by Engle and Rangel (2008) and Rangel and Engle (2009). Because idiosyncratic risk rises until 1997 while systematic risk falls, average correlation levels drop substantially during this period. The large increase in market volatility during the credit crisis at the end of the sample period is driven by a sharp increase in both idiosyncratic risk and average correlations.

4.5 Cross-Sectional Price of Variance and Correlation Risk

The results in the previous section show that market volatility, idiosyncratic risk, and correlations strongly fluctuate through time. Furthermore, the analysis reveals movements in these risk measures at different frequencies. Changes in high-frequency components of volatilities and correlations are transitory but movements in low-frequency components are persistent. In this section I examine whether these risk dynamics generate intertemporal hedging demands and significant risk premia in the cross-section of returns, as predicted by the ICAPM in (4.4.1). In particular, I define daily volatility and correlation factors as the first differences

Figure 4.6: Average Long and Short Run Idiosyncratic Volatility

This figure plots the value-weighted average of the high- and low-frequency components of the idiosyncratic volatility of 30 Dow Jones stocks for the period January 1990 to December 2008. The high-frequency component is modeled as a unit GARCH process and the low-frequency component is modeled as an exponential quadratic spline. The estimation is based on daily returns on the Dow Jones stocks and the resulting volatilities are annualized.

of the high- and low-frequency components of the conditional market volatility, value-weighted average idiosyncratic volatility, and average correlation obtained from the FSG-DCC model. I control for size, value, and momentum factors from the data library of Kenneth French.¹⁹ I also include a liquidity factor, which is calculated as the innovation in the value-weighted average of the Amihud (2002) illiquidity measure across all stocks in the CRSP universe.

Table 4.5 presents descriptive statistics for the daily volatility and correlation factors and for the other pricing factors. By construction, the innovations in the high-frequency components of volatilities and correlations are strongly time-varying while the shocks to the low-frequency parts are more stable, particularly at the daily frequency. Because correlations between the long and short run factors are low, they can capture orthogonal sources of risk. The table further shows that the high-(low)-frequency idiosyncratic volatility and correlation factors are positively correlated with the short (long) run market volatility factor, consistent with the decomposition of market volatility into individual volatilities and correlations. Interestingly, the innovation in low-frequency market volatility is more strongly correlated with the innovation in low-frequency idiosyncratic volatility whereas

¹⁹See http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

Figure 4.7: Average Long and Short Run Correlation

This figure plots the weighted average of the high- and low-frequency components of the correlations between 30 Dow Jones stocks for the period January 1990 to December 2008. These high- and low-frequency conditional correlations are obtained from the Factor-Spline-GARCH-DCC model described in Section 4.4.2.

the short run market volatility factor is more correlated with the short run correlation factor. Table 4.5 also indicates that the correlations between the volatility and correlation factors and the size, value, momentum, and liquidity factors are low. These factors can therefore have explanatory power for the cross-section of returns beyond the explanatory power of the traditional risk factors.²⁰

Importantly, the ICAPM in equation (4.4.1) specifies a relation between returns and *conditional* covariances. Ghysels, Santa-Clara, and Valkanov (2005) find a positive risk-return tradeoff at the market level when using a mixed-data sampling estimator of conditional variance. In contrast, Goyal and Santa-Clara (2003) find an insignificant relation when using an unconditional measure of market variance. Bali (2008) and Bali and Engle (2008) also show that the use of conditional covariances is crucial in identifying a positive risk-return tradeoff. Moreover, given the empirical evidence on time-varying betas (see, for instance, Bollerslev, Engle, and Wooldridge (1988) and Jagannathan and Wang (1996)), I allow for time variation in these conditional covariances.

²⁰Ang, Hodrick, Xing, and Zhang (2006) show that correlations between the market variance factor and the traditional pricing factors are somewhat higher when measured at a monthly frequency but are still low, except for the correlation with the liquidity factor. Nevertheless, they find that even after controlling for liquidity risk, market variance risk remains significantly priced in the cross-section.

Table 4.5: Summary Statistics for Risk Factors

This table presents descriptive statistics for daily variance and correlation factors and for traditional risk factors. R_M is the excess return on the S&P 500. $HF\ MVOL$ and $LF\ MVOL$ are innovations in high- and low-frequency market volatility, respectively. $HF\ IVOL$ and $LF\ IVOL$ are innovations in the value-weighted average of the high- and low-frequency idiosyncratic volatilities of the stocks in the Dow Jones. $HF\ COR$ and $LF\ COR$ are innovations in the weighted average of the high- and low-frequency correlations between the stocks in the Dow Jones. High- and low-frequency components of volatilities and correlations are obtained from the Factor-Spline-GARCH-DCC model outlined in Section 4.4.2 and are multiplied by 100 for ease of exposition. SMB , HML , and UMD are the daily size, value, and momentum factors from the data library of the Kenneth French. $ILLIQ$ is the innovation in the value-weighted average of the Amihud (2002) illiquidity measure across all stocks in the CRSP universe. The sample is daily and covers the period from January 1990 to December 2008.

	Mean	Std. Dev.	R_M	HF	LF	HF	LF	HF	LF	HF	LF	SMB	HML	UMD	ILLIQ
R_M	0.0188	1.13	1.00												
HF MVOL	0.0215	8.81	0.09	1.00											
LF MVOL	0.0127	0.11	-0.02	0.01	1.00										
HF IVOL	0.0291	3.56	0.04	0.18	0.04	1.00									
LF IVOL	0.0348	0.19	-0.02	0.00	0.86	0.03	1.00								
HF COR	0.0004	2.69	0.05	0.88	0.00	-0.05	0.00	1.00							
LF COR	-0.0020	0.02	0.00	0.00	0.37	0.00	-0.11	0.01	1.00						
SMB	0.0036	0.57	-0.25	-0.08	-0.01	-0.01	-0.01	-0.10	0.01	1.00					
HML	0.0161	0.58	-0.33	0.00	-0.01	0.00	-0.02	0.01	0.02	-0.16	1.00				
UMD	0.0458	0.82	-0.20	0.00	0.02	-0.04	0.01	0.00	0.00	0.09	-0.17	1.00			
ILLIQ	-0.0010	0.49	-0.09	-0.07	0.00	0.01	0.00	-0.03	0.00	-0.01	0.01	0.00	1.00		

Following Bali and Engle (2008), I estimate the time-varying conditional covariances between the stock returns and the risk factors using the standard DCC model of Engle (2002).²¹ Specifically, I first estimate univariate GARCH models for all returns and for the risk factors to obtain estimates of their conditional variances. I use these estimates to compute standardized residuals and estimate the conditional correlations between these residuals. The estimates of conditional variances and correlations are then used to construct the conditional covariances. Bali and Engle (2008) stress that time series and cross-sectional data should be pooled to gain statistical power and find a significant risk-return relation. Pooling is possible because the ICAPM implies that the intertemporal relation should be equal across assets. In the estimation I therefore impose the constraint that the γ and λ parameters in (4.4.1) should be the same across stocks. I estimate this system of equations using the seemingly unrelated regression approach (SUR). The SUR model accounts for heteroskedasticity, autocorrelation, and contemporaneous correlation across residuals to estimate the system more efficiently by FGLS.²² Details are in Bali (2008).

I test whether long and short run market volatility risk, idiosyncratic volatility risk, and correlation risk are priced by examining the estimates of the elements in the vector λ , reported in Table 4.6. The first column corresponds to the traditional risk-return tradeoff without intertemporal hedging demand. The ICAPM implies that the estimated γ should be equal to the coefficient of relative risk aversion. The coefficient estimate is 2.86 with a t -statistic of 3.76, which is economically plausible and in line with the findings of Bali and Engle (2008). I also test the prediction of the ICAPM that all intercepts are equal to zero. The last row in column 1 shows that the Wald statistic for testing this joint hypothesis is 15.70, with a p -value of 0.99. Hence, I cannot reject the hypothesis that all pricing errors are zero, which is evidence in favor of the conditional ICAPM.

Column two shows that the price of long run market volatility risk is significantly negative. The estimates in column three indicate that high-frequency market volatility risk is also negatively priced. Column four reports results when the long and short run market volatility factors are simultaneously included in the regression. Both factors remain significant and their prices of risk are almost the same as in the specifications in which they are included separately. Moreover, the

²¹The conditional covariances between the stock returns and the market factor can be obtained directly from the FSG-DCC model.

²²A common concern in asset pricing tests is estimation error in factor exposures, which are the conditional covariances in this chapter. Since these covariances are obtained from a DCC model, standard methods to account for measurement error like the Shanken (1992) correction cannot be applied. Bali and Engle (2008) assess the impact of measurement error by comparing results of a one-step multivariate GARCH-in-mean estimation, in which the conditional covariances and prices of risk are estimated simultaneously, to results obtained from the two-step procedure used in this chapter, in which these components are estimated sequentially. The one-step estimation is computationally very intensive and can only be used to test whether the market factor is priced in the ICAPM. Bali and Engle (2008) find that the one-step and two-step methods provide similar evidence on the relation between the risk and return on Dow stocks and conclude that measurement errors in conditional covariances do not have significant effects on the magnitude and statistical significance of the risk aversion coefficient.

market factor remains significant after including the two volatility factors. The negative prices of market volatility risk imply that stocks that covary positively with unexpected changes in long and short run aggregate volatility have a positive hedging demand and therefore higher prices and lower expected returns. These findings extend the results of Adrian and Rosenberg (2008), who show that innovations in long- and short-term market volatility carry a negative price of risk in a cross-section of size-B/M portfolios, to a sample of individual stocks.

I examine the pricing of innovations in high- and low-frequency components of average idiosyncratic volatility in columns five and six. Long run idiosyncratic volatility carries a negative price of risk, which is significant at the 10% level. In contrast, short run idiosyncratic volatility risk is not priced in the cross-section of Dow stocks. This is consistent with the low correlation between the short run components of market volatility and idiosyncratic volatility in Table 4.5 and implies that short run market volatility risk is not priced because of short-term idiosyncratic volatility risk. These results remain unchanged when the two idiosyncratic volatility factors enter the model simultaneously.

The price of long run correlation risk in column seven is negative and significant at a 1% level. Short run correlation risk is also negatively priced and significant at a 5% level. When both correlation factors are jointly added to the regression, their coefficients and *t*-statistics do not change much, consistent with the low correlation between these two factors in Table 4.5. This indicates that the long and short run correlation factors capture orthogonal sources of risk. The last three columns in Table 4.6 show that when innovations in high- and low-frequency idiosyncratic volatilities and correlations are included together, the long run components dominate in terms of statistical significance. Short run correlation risk is still significant at a 10% level but innovations in short-term idiosyncratic risk are insignificant.

In Table 4.7 I study the pricing of long and short run volatility and correlation risk after controlling for size, value, momentum, and liquidity factors. I first estimate the price of these factors excluding the volatility and correlation factors. The first two columns show that *SMB* and *ILLIQ* have the expected signs but are insignificant, which is not surprising given that the 30 Dow stocks used as test assets are all large and liquid. *HML* is positively priced but loses its significance after including the momentum factor, which is significant at the 5% level. Columns three and four show that the long and short run market volatility factors remain priced after controlling for the standard risk factors. The innovation in average long run idiosyncratic volatility also continues to carry a significantly negative price of risk at the 10% level after including the other factors. The estimates of the prices of long and short run correlation risk are also very similar to those reported in Table 4.6. The finding that the pricing of variance and correlation risk is unaffected by the inclusion of the other factors suggests that the variance and correlation factors capture risks that are different from the traditional size, value, momentum, and liquidity risks.

My finding that innovations in long run idiosyncratic risk and long-term correlations are priced factors suggests that long run market volatility risk is priced

Table 4.6: Cross-Sectional Price of Volatility and Correlation Risk

This table reports estimates of the cross-sectional price of volatility and correlation risk. The intertemporal risk-return tradeoff is estimated as a system of equations for a panel of 30 Dow Jones stocks using the SUR approach. The dependent variables are the daily excess stock returns and the independent variables are the conditional covariances between the stock returns and risk factors. These factors are the daily excess market return and the high- and low-frequency volatility and correlation factors defined in Table 4.5. The coefficients for the high- and low-frequency correlation factors are divided by 10 for ease of exposition. The sample period is from January 1990 to December 2008. *t*-Statistics adjusted for heteroskedasticity, autocorrelation, and cross-correlations among the residuals are in parentheses. *Wald* is the Wald test statistic for the null hypothesis that all intercepts are jointly equal to zero.

R_M	2.86 (3.76)	2.30 (2.83)	3.58 (4.58)	3.05 (3.66)	2.49 (3.13)	2.82 (3.68)	2.41 (3.01)	3.11 (4.07)	3.09 (4.02)	3.31 (4.28)	2.34 (2.94)	3.05 (3.94)	2.43 (3.00)
HF MVOL	-1.08 (-4.08)	-1.06 (-4.01)											
LF MVOL		-89.60 (-2.00)	-82.72 (-1.85)										
HF IVOL				0.56 (0.75)			0.70 (0.93)				0.52 (0.69)		0.88 (1.16)
LF IVOL					-28.67 (-1.67)	-30.25 (-1.75)						-63.81 (-3.49)	-66.36 (-3.60)
HF COR							-0.60 (-1.95)			-0.53 (-1.72)	-0.59 (-1.92)		-0.52 (-1.69)
LF COR								-125.41 (-2.67)	-123.21 (-2.58)			-159.38 (-3.57)	-158.92 (-3.55)
Wald	15.70 (0.99)	18.42 (0.95)	17.00 (0.97)	19.70 (0.92)	15.70 (0.99)	16.86 (0.97)	16.92 (0.97)	18.29 (0.95)	22.70 (0.83)	24.07 (0.77)	18.15 (0.96)	29.35 (0.50)	31.07 (0.41)
<i>p</i> -value													

Table 4.7: Price of Volatility and Correlation Risk Controlling for Traditional Risk Factors

This table reports estimates of the cross-sectional price of volatility and correlation risk for a panel of 30 Dow Jones stocks. The dependent variables are the daily excess stock returns and the independent variables are the conditional covariances between the stock returns and risk factors. These factors are the daily excess market return, the high- and low-frequency volatility and correlation factors, and the size, value, momentum, and illiquidity factors defined in Table 4.5. The coefficients for the high- and low-frequency correlation factors are divided by 10 for ease of exposition. The sample period is from January 1990 to December 2008. *t*-Statistics adjusted for heteroskedasticity, autocorrelation, and cross-correlations among the residuals are in parentheses.

R_M	2.68 (3.25)	2.03 (2.08)	2.89 (3.20)	2.14 (2.15)	2.31 (2.65)	1.80 (1.74)	3.25 (3.90)	2.14 (2.20)	2.44 (2.78)	1.85 (1.88)
HF MVOL			-1.06 (-3.98)	-1.05 (-3.92)						
LF MVOL			-82.22 (-1.99)	-74.50 (-1.87)						
HF IVOL					0.70 (0.92)	0.71 (0.93)			0.86 (1.14)	0.87 (1.14)
LF IVOL					-29.83 (-1.82)	-27.19 (-1.72)			-55.78 (-2.87)	-46.68 (-2.25)
HF COR							-0.54 (-1.73)	-0.53 (-1.70)	-0.53 (-1.71)	-0.55 (-1.72)
LF COR							-133.36 (-2.89)	-132.95 (-2.86)	-159.46 (-3.54)	-154.39 (-3.32)
SMB	1.40 (0.62)	2.08 (0.84)	2.16 (0.92)	2.38 (1.00)	1.73 (0.73)	2.10 (0.88)	0.79 (0.30)	1.30 (0.54)	1.06 (0.44)	1.33 (0.55)
HML	4.32 (2.15)	1.57 (0.67)	2.84 (1.32)	1.20 (0.50)	3.51 (1.64)	1.45 (0.61)	5.74 (2.82)	3.04 (1.28)	3.84 (1.79)	2.49 (1.04)
UMD		-4.29 (-2.00)		-3.06 (-1.33)		-3.91 (-1.71)		-4.30 (-1.99)		-2.51 (-1.09)
ILLIQ		-5.66 (-0.86)		-6.96 (-1.06)		-6.37 (-0.96)		-4.44 (-0.67)		-4.22 (-0.63)

because of both long run idiosyncratic volatility risk and correlation risk. In contrast, short run market volatility risk is only priced in the cross-section because of the negative price of short-term correlation risk. To study the interaction between market volatility risk, idiosyncratic volatility risk, and correlation risk in more detail and to assess the economic magnitude of these risks, I calculate the risk premia of each stock on the high- and low-frequency volatility and correlation factors. These risk premia are computed by multiplying the factor exposures (i.e., the conditional covariances) by the corresponding prices of market volatility risk in column four of Table 4.6 and the prices of idiosyncratic volatility risk and correlation risk in the last column of that table. For ease of interpretation I aggregate the daily risk premia to a monthly frequency.

The first column in Table 4.8 reveals a wide spread in the premium for short run market volatility risk, ranging from -0.37% per month for Intel to 0.42% for Alcoa. In general, most growth stocks have a positive exposure to short-term market volatility risk, whereas most value stocks load negatively on this factor. Since the price of short run market volatility risk is negative, these loadings imply that for growth stocks the risk premium associated with short run market volatility is negative, whereas for value stocks it is positive. A similar but less pronounced pattern can be observed in the risk premium for long run market volatility risk, which has a cross-sectional standard deviation of 0.10%. Thus, it appears that growth stocks earn lower returns than value stocks because they can be used to hedge against sudden changes in market volatility. Adrian and Rosenberg (2008) argue that the positive loadings of growth stocks on the market volatility factors could be driven by investor learning about the growth opportunities of these firms (see Pastor and Veronesi (2003)). Alternatively, they suggest that the cross-sectional differences in exposures to volatility risk may reflect differences in the option value of growth opportunities that are due to heterogeneous adjustment costs, as documented by Zhang (2005).

Column four reports the risk premium for the high-frequency component of idiosyncratic volatility risk, which has a much narrower spread than the other risk premia. This is in line with the insignificance of the short run idiosyncratic volatility risk factor in the ICAPM regressions in Tables 4.6 and 4.7. In contrast, the long run idiosyncratic volatility risk premium varies substantially across stocks. Similar to the premium for market volatility risk, it is negative for most growth stocks and positive for the majority of value stocks.

The premium for short run correlation risk has the largest cross-sectional standard deviation of all premia in Table 4.8. It ranges from -0.41 for Intel to 0.43 for Alcoa. Strikingly, this premium lines up very well with the premium for high-frequency market volatility risk. The last column shows that the cross-sectional dispersion in the premium for long-term correlation risk is similar to the variation in the premium for low-frequency market volatility risk.

Panel B summarizes the similarity in variance and correlation premia by showing their cross-sectional correlation. The premium for short-term market volatility risk is positively related to the premium for high-frequency correlation risk but negatively correlated with the premium for short run idiosyncratic volatility risk.

This finding provides further evidence that short-term market volatility risk is priced in the cross-section of returns because short run correlation risk is priced. In contrast, the premium for long run market volatility risk is positively associated with the premia for both long run idiosyncratic volatility risk and long run correlation risk.

Overall, the results in this section show that shocks to high- and low-frequency components of market volatility, average idiosyncratic volatility, and market-wide correlations have significant prices of risk in the cross-section of individual stock returns. The variance and correlation factors continue to have explanatory power after controlling for traditional size, value, momentum, and liquidity factors, which justifies their role as separate risk factors in the ICAPM.

4.6 Conclusion

This chapter examines the pricing of long and short run market volatility risk in aggregate and individual returns by decomposing market volatility risk into correlation risk and idiosyncratic volatility risk. At the aggregate level, I show that the predictive power of the market variance premium for market returns documented by Bollerslev, Tauchen, and Zhou (2009) is completely driven by the correlation risk premium. In contrast, the average individual variance risk premium does not predict the equity premium. The predictive power of the correlation risk premium is robust to the inclusion of traditional return predictors and is even larger than that of the market variance risk premium for one-month ahead returns.

I investigate the cross-sectional pricing of long and short run variance and correlation risk by decomposing market volatility, idiosyncratic volatility, and correlations into high- and low-frequency components using a semi-parametric approach proposed by Rangel and Engle (2009). I identify distinct cyclical patterns in long-term volatilities and correlations that highlight the importance of allowing unconditional volatilities and correlations to vary over time. The intertemporal CAPM predicts that these risk dynamics are priced in the cross-section because investors want to hedge against a sudden increase in volatilities and correlations.

I find that innovations in the long-term component of market volatility are negatively priced in the cross-section because both shocks to long run idiosyncratic volatility and shocks to long-term correlations carry a negative price. Short run market volatility is priced because of the negative price of short-term correlation risk. The cross-sectional explanatory power of the variance and correlation risk factors is not absorbed by size, value, momentum, and liquidity factors. Furthermore, I show that for most growth stocks the variance and correlation risk premia are negative, since these stocks provide insurance against shocks to volatilities and correlations. In contrast, value firms have higher expected returns to compensate for their poor performance when uncertainty increases and diversification benefits decrease. Future research should seek to identify the underlying economic sources of long and short run variance and correlation risk and explain the differences in the exposures to these risks across firms.

Table 4.8: **Factor Risk Premia**

This table reports risk premia of 30 DJIA stocks on volatility and correlation factors. The premia are computed by multiplying factor exposures by the corresponding prices of volatility risk and correlation risk in Table 4.6. The daily risk premia are aggregated to a monthly frequency.

	HF MVOL	LF MVOL	HF IVOL	LF IVOL	HF COR	LF COR
<i>Panel A: Factor Risk Premia</i>						
AA	0.42	0.02	-0.02	0.11	0.43	-0.04
AIG	0.38	0.33	-0.05	0.59	0.29	0.18
AXP	0.29	0.01	-0.04	0.09	0.20	-0.08
BA	0.19	0.14	0.01	0.26	0.16	0.12
BAC	0.25	0.21	0.06	0.35	0.23	0.06
C	0.33	0.18	0.08	0.18	0.33	0.09
CAT	0.33	0.12	0.05	0.22	0.35	-0.01
CVX	0.09	0.00	-0.03	-0.02	-0.03	0.02
DD	0.21	0.03	0.02	0.00	0.17	0.00
DIS	0.13	0.04	0.04	0.00	0.06	0.06
GE	-0.04	0.01	0.05	-0.01	-0.07	0.02
GM	0.26	0.14	0.11	0.22	0.27	0.00
HD	0.22	0.09	0.03	-0.20	0.17	0.24
HPQ	0.02	-0.10	0.13	-0.09	0.01	-0.12
IBM	-0.10	-0.16	0.05	-0.15	-0.20	-0.10
INTC	-0.37	-0.04	0.10	-0.05	-0.41	0.02
JNJ	0.08	0.00	-0.02	-0.03	-0.04	0.05
JPM	0.24	-0.04	0.05	-0.10	0.21	-0.09
KO	-0.01	0.03	0.01	-0.02	-0.11	0.11
MCD	0.26	-0.08	-0.01	-0.15	0.24	0.00
MMM	0.26	0.08	0.00	0.11	0.19	0.03
MRK	0.03	-0.07	0.01	-0.17	-0.04	-0.06
MSFT	-0.08	0.01	0.09	-0.18	-0.12	0.17
PFE	0.29	-0.01	-0.07	-0.06	0.20	0.08
PG	-0.06	0.04	0.02	0.03	-0.16	0.00
T	0.15	-0.06	0.03	-0.04	0.04	-0.05
UTX	0.41	0.00	0.01	0.12	0.40	-0.03
VZ	-0.02	-0.08	0.10	-0.10	-0.08	-0.06
WMT	-0.05	-0.08	0.07	-0.30	-0.09	0.07
XOM	-0.11	0.01	0.02	-0.01	-0.22	0.03
<i>Panel B: Cross-Sectional Correlation of Factor Risk Premia</i>						
HF MVOL	1.00					
LF MVOL	0.51	1.00				
HF IVOL	-0.44	-0.17	1.00			
LF IVOL	0.54	0.86	-0.22	1.00		
HF COR	0.98	0.51	-0.30	0.53	1.00	
LF COR	0.07	0.60	-0.17	0.20	0.05	1.00

Chapter 5

Option Trading and Individual Investor Performance*

This chapter examines the impact of option trading on individual investor performance. The results show that most investors incur substantial losses on their option investments, which are much larger than the losses from equity trading. We attribute the detrimental impact of option trading on investor performance to poor market timing that results from overreaction to past stock market returns. High trading costs further contribute to the poor returns on option investments. Gambling and entertainment appear to be the most important motivations for trading options while hedging motives only play a minor role. We also provide strong evidence of performance persistence among option traders.

5.1 Introduction

Over the last decade, Internet brokerage has dramatically changed the investment landscape. The professional traders who used to dominate financial markets now find themselves surrounded by a much larger and more divergent crowd: individual investors. To gain a better insight into the trading behavior of these individual investors, financial economists examine their performance using trading records and position statements obtained from brokerage firms.¹

A growing literature presents evidence of irrational behavior of individual investors in option markets. Poteshman (2001) shows that option market investors

*This chapter is based on Bauer, Cosemans, and Eichholtz (2009), published in *Journal of Banking and Finance*.

¹Studies have considered the trading behavior and performance of individual investors in, among others, the United States (Barber and Odean (2000)), China (Ng and Wu (2007)), Israel (Shapira and Venezia (2001)), and Sweden (Anderson (2008)).

exhibit the same pattern of short-term underreaction and long-term overreaction to information that has been found in stock markets. In addition, Poteshman and Serbin (2003) find that customers of discount brokers regularly engage in irrational early exercise of stock options. Mahani and Poteshman (2008) document that discount clients act irrationally by entering option positions that load up on growth stocks a few days before earnings announcements, even though at earnings announcements value stocks usually outperform substantially. Moreover, Han (2008) finds that investor sentiment about the stock market affects option prices. Another strand of literature stresses the importance of irrational determinants of individual investors' trading activity in stock markets, like gambling (Kumar (2009b)), entertainment (Dorn and Sengmueller (2009)), and sensation-seeking (Grinblatt and Keloharju (2009)). However, options seem to be more attractive for these purposes than stocks due to the leverage they provide and the positive skewness of their payoffs. Indeed, Lakonishok, Lee, Pearson, and Poteshman (2007) show that a large fraction of individuals' option activity is motivated by speculation on the direction of future stock price movements.

We extend this work by studying the impact of option trading on individual investor performance. Doing so gives us the opportunity to shed light on the question whether individual investors understand the risk and return characteristics of these more complex securities and whether they are able to use these instruments successfully. Previous work (e.g., Barber and Odean (2000)) has shown that excessive trading by individual investors leads to substantial losses on their common stock investments. We therefore examine both the absolute returns of option traders and their performance relative to those who only trade stocks. In addition, we identify the determinants of option trading volume at the investor level to examine whether option trading is related to investor characteristics that have been linked to gambling and sensation seeking in stock markets by Kumar (2009b) and Grinblatt and Keloharju (2009). Furthermore, we investigate whether there is a group of option traders who are able to consistently earn abnormal returns. This is motivated by the findings of Coval, Hirshleifer, and Shumway (2005), who present evidence of performance persistence among a small group of stock investors. We then analyze the characteristics and trading strategies of these skilled option investors.

We perform the empirical analysis using a unique database that comprises more than 68,000 accounts and more than eight million trades in stocks and options at a large online broker in the Netherlands. In terms of size, the sample is comparable to the data set often used in studies for the United States (see, e.g., Barber and Odean (2000) and Kumar (2009b)). We examine investor behavior and performance from January 2000 to March 2006, which covers the height of the stock market boom in 2000, the subsequent bust in stock prices in 2001 and 2002, and the recovery from 2003 to 2006. Thus, we are able to examine whether major market movements affect trading behavior and investor performance.

We use several methods to deal with the specific risk and return characteristics of individual investor portfolios. First, to adjust returns for risk and style tilts, we use a multifactor model in the spirit of Agarwal and Naik (2004) to capture

the nonlinear payoffs of options. Furthermore, we use a Kalman filter approach to allow for time variation in risk loadings and style preferences. Finally, we introduce an approach that allows us to control for risk and style exposures even when the number of time series observations for some investors is very small.

We find that option traders incur much larger losses on their investments than equity traders. The gross return difference between these two groups of investors equals more than 1% a month, after taking risk and style differences into account. Controlling for known determinants of investor returns like gender, age, turnover, account value, income, and experience does not explain the return differential between option investors and equity investors. Instead, we attribute the poor performance of option traders to bad market timing that results from overreaction to past stock market movements. We construct a call/put ratio based on the option trades of the clients of the broker and find this ratio to be highly correlated with two other sentiment indicators, the consumer confidence index and the VIX index. In addition, estimation results for a vector autoregressive model show that the call/put ratio is driven by past market returns.

We further document that the demographic (age and gender), socioeconomic (income), and portfolio characteristics (account value and turnover) of option traders and equity traders are very similar. In particular, the empirical results show that men are more likely than women to engage in both option trading and equity trading and also exhibit a higher trading intensity. Furthermore, we extend the result of Anderson (2008) that poorer investors trade most to the option market. However, an important difference between option traders and equity traders is the impact of past portfolio returns. Specifically, while past returns do not affect an individual's trading activity in stock markets, they have a significantly negative influence on option trading volume. This is related to the finding of Coval and Shumway (2005) that futures traders are highly loss-averse and take more risk following losses than following gains.

We link the option positions that investors take to their common stock holdings and show that most investors do not use options for hedging underlying stock positions, which confirms the results of Lakonishok, Lee, Pearson, and Poteshman (2007) for the U.S. market. Instead, our finding that single men with low income and little investment experience trade most suggests that gambling and sensation-seeking are important determinants of option trading, consistent with the results of Kumar (2009b) and Grinblatt and Keloharju (2009) for equity investors. This is confirmed by the responses of a group of brokerage clients to several statements on investment attitude, which reveal that the majority of investors enjoy trading and only invest the money they do not directly need. The responses also show that only a small subset of investors uses the option Greeks when trading options. Hence, a lack of knowledge about the risk and return characteristics and the use of options might be another explanation for the detrimental impact of option trading on investor performance. Seru, Shumway, and Stoffman (2009) provide support for this hypothesis by showing that inexperienced investors learn slowly.

Consistent with this interpretation, we identify a small group of sophisticated option traders who succeed in consistently outperforming other investors. Specif-

ically, option traders who are in the top decile portfolio based on past one-year performance continue to outperform investors in the bottom decile over the next year in terms of both gross and net alphas. Performance persistence is somewhat weaker on shorter horizons but still significant for six-month periods. Persistence in trading costs explains only part of total performance persistence, as we also find persistence in gross performance. Analyzing the composition of the decile portfolios shows that investors in the bottom deciles tend to hold small accounts with high turnover. Furthermore, these accounts are predominantly held by men with low income and little investment experience.

The chapter proceeds as follows. In Section 5.2 we give a short overview of individual investors and financial markets in the Netherlands and introduce the data set of investor accounts and trades that is used in the empirical analysis. Section 5.3 outlines the methods used for performance calculation and attribution and Section 5.4 presents the empirical results. Section 5.5 concludes.

5.2 Data

5.2.1 Individual Investors and Financial Markets in the Netherlands

We use a data set of individual investor accounts at a large online discount broker in the Netherlands. At the end of 2006, 224 companies had listed their shares at Euronext Amsterdam, of which 128 were domestic firms and 96 were foreign.² The total market capitalization of these companies was 591 billion euros. Total equity turnover in 2006 was close to 687 billion euros of which 585 billion euros was turnover of stocks included in the AEX index. The Dutch AEX stock market index is a value-weighted index of the 25 most actively traded firms with a combined market value of 50 billion euros. On the Euronext Liffe Amsterdam exchange, options on approximately 50 different stocks are traded. In 2006 a total of 89 million stock option contracts and 25 million index option contracts were traded.

An annual survey performed by the marketing research agency Millward-Brown in 2006 among approximately 1,000 Dutch investors characterizes the market of individual investors in the Netherlands. In total, 1.5 million out of 7 million Dutch households invest their money in financial markets. The average Dutch retail investor trades securities at 1.4 different financial institutions (banks and online brokers). Thus, it is unlikely that the investors in our sample trade at many other institutions than the online broker. Almost half of Dutch retail investors trade through websites of banks and Internet brokers. These online investors trade almost three times as much as offline investors, in line with results documented by Barber and Odean (2002). The value-weighted asset mix of the average Dutch investor consists of 44% mutual funds, 34% stocks, 15% bonds, 5% derivatives, and 2% other securities. Investors indicate that they predominantly invest in mutual

²These statistics are retrieved from the Euronext Global factbook, available at <http://www.euronext.com>.

funds via a traditional bank. In contrast, online brokers are mostly used for trading stocks and derivatives. For the average (median) investor the total value of the investment portfolio is €60,000 (€35,000).

5.2.2 Data Set

The raw data set contains the daily trades and end-of-the-month portfolio positions of all individual investor accounts that existed during the period from January 2000 to March 2006. Due to trading restrictions, we exclude accounts owned by minors (age < 18 years). We further exclude accounts with a beginning-of-the-month value of less than €250 to reduce the impact of outliers. Imposing these restrictions leaves 68,146 accounts and more than two million monthly portfolio overviews. On average, investors are present in the sample for 36 months.

Table 5.1 presents summary statistics for the data set. The sample is split into 26,266 option traders and 41,880 equity traders. We define an option trader as an investor who trades options at least once during the sample period. The proportion of men among option investors and equity investors is similar (more than 75%). Thus, if we take gender as a proxy for overconfidence, as suggested by Barber and Odean (2001), option traders and equity traders exhibit the same level of overconfidence. The average age of option traders and equity investors is also very similar (45 years). For option investors the mean (median) number of trades during the sample period is 80 (16) and for those who only invest in stocks the average (median) number of trades is 45 (11). The ten most traded stocks account for 34% of all stock transactions and almost half of the value of all stock trades. The four most frequently traded stocks all belong to the IT and Telecommunications sector, which reflects the heavy trading of these stocks during the tech bubble. Most other frequently traded stocks are those of large companies.³ A striking feature of our sample is that almost half of all trades are in options.⁴ More than 33 million option contracts are traded, which, based on historical statistics retrieved from NYSE Euronext Liffe, is 7.2% of all option contracts traded at the Amsterdam options exchange during the sample period. The 10 most traded underlyings account for 77.2% of all option trades and 72.0% of the value of all option trades. Options on the Dutch AEX index account for 47.8% of all option transactions and 34.3% of the value of all option trades. The most frequently traded stock options are those on the largest firms with the most liquid stocks.⁵

³By constructing value-weighted factors for performance attribution we take into account that investors mainly hold large caps, since these stocks receive more weight in the construction of the factors.

⁴To make a clean comparison between the impact of option trading and equity trading on investor performance we disregard transactions in bonds and futures contracts, which constitute only a small fraction of all trades.

⁵The option-based factors that we include in the performance evaluation model are constructed using AEX index options. Large and liquid stocks receive more weight in the value-weighted AEX index. Therefore, the fact that options are not available for small and illiquid stocks does not affect the performance evaluation.

Table 5.1 also reports statistics for monthly turnover per account, which we define as the average of the value of all security purchases and sales divided by beginning-of-the-month account value. Average turnover for option traders is 8.9% but median turnover is zero. For equity traders the average (median) turnover is 23.7% (zero). These results show that although the majority of investors (65%) do not trade on a monthly basis, a subset of investors trades very often. The fact that for option investors the average number of trades is higher than for equity investors while the average turnover is lower, indicates that the value of each option transaction is smaller than the value of a stock transaction. Consequently, the average costs per transaction are higher for option traders than for equity traders (€32 and €22, respectively).

The average (median) portfolio size of option traders and equity traders is again close to each other, with a mean value of €35,000 and a median of €10,000. Combining the average account value with a total portfolio value of €60,000 for the average Dutch investor, as reported by Millward-Brown (2006), shows that the average client invests more than half of the total investment portfolio at the online broker. We also assign income to investors, based on their zip-code and income data retrieved from the Dutch Central Bureau of Statistics (CBS). The average gross income of both groups of traders is slightly higher than €2,500 a month. Following Seru, Shumway, and Stoffman (2009), we measure experience by the number of months that an investor has been trading. The results in Table 5.1 show that on average, option traders are a bit more experienced than equity traders (35 and 33 months of experience, respectively). In general, however, the descriptive statistics show a striking similarity between the demographic, socioeconomic, and portfolio characteristics of option traders and equity traders.

5.3 Methods

5.3.1 Measuring Investor Performance

We define investor performance as the monthly change in the market value of all stocks and options in an investor's account. End-of-the-month account value is net of transaction costs the investor incurred during the month. Since we measure performance on a monthly basis, we have to make an assumption concerning the timing of deposits and withdrawals of cash and securities. To be conservative, we assume that deposits are made at the beginning of the month and that withdrawals take place at the end of the month.

We also performed the analysis under the assumption that deposits and withdrawals are made halfway the month and find that our results are robust to this assumption. Thus, we calculate net performance as

$$R_{it}^{net} = \frac{(V_{it} - V_{it-1} - NDW_{it})}{(V_{it-1} + D_{it})}, \quad (5.3.1)$$

where V_{it} is the account value at the end of month t , NDW_{it} is the net of deposits and withdrawals during month t , and D_{it} are the deposits made during month t .

Table 5.1: Descriptive Statistics on Investor Accounts and Trades

This table presents descriptive statistics for a sample of 68,146 investor accounts at a Dutch online broker. We split the sample into 26,266 option traders and 41,880 equity traders. The sample period is from January 2000 to March 2006. The variables are defined as follows: *Gender* and *Age* are the gender and age of the primary accountholder. *Trades* is the total number of trades per account during the sample period. *Turnover* is the average of the value of all purchases and sales in a given month divided by the beginning-of-the-month account value. *Portfolio value* is the average market value of all assets in the investor's portfolio. *Income* is the average monthly income assigned to investors based on their zip-code. *Experience* is the number of months an investor has been trading.

	Mean	Std. Dev.	5th	25th	Median	75th	95th
Gender (male=1)							
Option traders	0.78						
Equity traders	0.76						
Age (years)							
Option traders	45.28	12.98	26	36	44	55	67
Equity traders	44.72	13.29	25	36	44	54	66
Trades (#)							
Option traders	80.32	304.23	1	4	16	59	331
Equity traders	45.26	135.43	1	4	11	35	185
Turnover (%)							
Option traders	8.85	100.83	0	0	0	0	25.60
Equity traders	23.68	184.96	0	0	0	2.38	83.33
Portfolio value (€)							
Option traders	34,682	111,344	500	2,700	9,419	29,432	133,655
Equity traders	35,053	179,762	662	3,343	10,004	28,351	129,812
Income (€)							
Option traders	2,548	935	1,500	1,900	2,300	2,900	4,300
Equity traders	2,563	937	1,500	2,000	2,400	2,900	4,400
Experience (months)							
Option traders	34.64	28.02	3.00	12.00	27.28	50.80	75.00
Equity traders	32.94	25.34	2.00	12.00	27.52	48.80	74.00

We obtain gross returns by adding back transaction costs incurred during month t , TC_{it} , to end-of-the-month account value,

$$R_{it}^{gross} = \frac{(V_{it} - V_{it-1} - NDW_{it} + TC_{it})}{(V_{it-1} + D_{it})}. \quad (5.3.2)$$

We only consider direct transaction costs (commissions) and do not add back any indirect transaction costs (market impact and bid-ask spreads). The trades of most individual investors are relatively small, so their market impact is likely to be limited. In addition, Keim and Madhavan (1998) point out that quoted bid-ask spreads may be imprecise estimates of the true spread, because trades are often executed inside the quoted spread. Therefore, Barber and Odean (2000) estimate the bid-ask spread using transaction prices and closing prices. However, this approach is inappropriate for our purposes as the resulting estimate of the spread includes the return on the trading day, which can be substantial in the case of options.

5.3.2 Performance Attribution

We attribute the returns on investor portfolios to different risk and style factors to obtain the abnormal performance. The Carhart (1997) four-factor model is used to adjust investor returns for exposures to the market, size, book-to-market, and momentum factors. We construct these factors for the Dutch market, since the investors in the sample mainly invest in Dutch securities.⁶ To characterize the market risk of the equity component of the portfolio returns, we include the return on the MSCI Netherlands equity index. We construct the factor-mimicking portfolios *SMB*, *HML*, and *MOM* according to the procedure outlined by Kenneth French.⁷

To characterize the nonlinear exposure from options, we build on the theoretical framework developed by Glosten and Jagannathan (1994), who propose adding option-based factors to performance attribution models. Agarwal and Naik (2004) implement this approach to measure the risk exposure of hedge funds and find that many funds use strategies that result in option-like payoffs. Following these studies, we include the excess returns on liquid at-the-money (ATM) European call and put options on the Dutch AEX stock market index to capture the nonlinear systematic risk exposure of investors' portfolios. In particular, at the end of each month, ATM call and put index options that expire two months later are bought. These index options are sold one month later and new index options are purchased. As a proxy for ATM options we select options whose strike prices are closest to the current index value. This rolling strategy of buying and selling index calls and

⁶In terms of volume (value) 95% (85%) of all trades are transactions in Dutch securities. This suggests the presence of a home bias among Dutch investors, which has previously been documented by French and Poterba (1991) for the U.S., Japan, and UK and by Karlsson and Norden (2007) for Sweden. Consequently, we find that Dutch versions of the factor-mimicking portfolios lead to a better model fit than do international factors.

⁷See http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

puts produces a time series of monthly returns on ATM calls and puts that we add to the performance attribution model.

We further include the value-weighted average return on Dutch stocks from the MSCI IT and Telecommunications sector to capture possible tech-related style tilts, since many economists, including Shiller (2005), argue that the technology bubble was fed by irrational euphoria among individual investors. Because the option-based factors and the IT factor are highly correlated with the market return, we orthogonalize them with respect to the market factor.

The time series model we estimate to obtain risk- and style-adjusted returns is

$$R_{it} = \alpha_i + \sum_{k=1}^K \beta_{ik} F_{kt} + \epsilon_{it}, \quad (5.3.3)$$

where R_{it} is the excess return on the portfolio of investor i , β_{ik} is the loading of portfolio i on factor k , and F_{kt} is the month t excess return on the k 'th factor-mimicking portfolio. The intercept α_i measures abnormal performance relative to the risk and style factors. The factor loadings indicate whether a portfolio is tilted towards a particular investment style.

Other studies on individual investor performance assume that factor loadings remain constant over time, i.e., they use unconditional or static models for performance attribution. However, a large body of empirical evidence shows that the systematic risk of stocks varies substantially over time as a function of the business cycle (see, e.g., Ferson and Harvey (1999) and Bauer, Cosemans, and Schotman (2009)). Moreover, in a dynamic world it is unlikely that investors keep their exposure to risk and style factors constant over time. Empirical support for this conjecture is provided by Kumar (2009a), who finds that individual investors exhibit time-varying style preferences, driven by past style returns and earnings differentials. Hence, Ferson and Schadt (1996) argue that fluctuations in factor exposures should be taken into account when measuring portfolio performance. We treat the time-varying alphas and betas as latent state variables and infer them directly from portfolio returns using the Kalman filter. We assume a random walk process for the latent alphas and betas.⁸ Specifically, we consider the following state-space representation:

$$R_{it} = \alpha_{it} + \sum_{k=1}^K \beta_{ikt} F_{kt} + \epsilon_{it}, \quad \epsilon_{it} \sim N(0, \sigma_{\epsilon}^2), \quad (5.3.4)$$

$$\alpha_{it} = \alpha_{it-1} + \nu_{it}, \quad \nu_{it} \sim N(0, \sigma_{\nu}^2), \quad (5.3.5)$$

$$\beta_{it} = \beta_{it-1} + \eta_{it}, \quad \eta_{it} \sim N(0, \Omega), \quad (5.3.6)$$

where ϵ_{it} , ν_{it} , and η_{it} are normally distributed mean zero shocks orthogonal to each other and with variance σ_{ϵ}^2 , σ_{ν}^2 , and diagonal covariance matrix Ω , respectively.

⁸We have also experimented with other specifications for the time-varying factor loadings, including mean-reverting processes. Because we find that most betas are very persistent and the number of time series observations available for estimation is limited, we choose the random walk specification, which is more parsimonious.

We test the null hypothesis of constant betas, which corresponds to the restriction that the diagonal elements of Ω are zero, using a likelihood ratio test. Equation (5.3.4) is the observation equation and equations (5.3.5) and (5.3.6) are the state equations. The Kalman filter is a recursive algorithm for sequentially updating the one-step-ahead estimate of the state mean and variance. We use it to calculate maximum likelihood estimates of the model parameters σ_{ic}^2 , σ_{iv}^2 , and Ω along with minimum mean-square error estimates of the state variables α_{it} and β_{it} . We set the initial one-step-ahead predicted values for the states equal to the OLS estimates from the static model. We treat the initial one-step-ahead predicted values of σ_{iv}^2 and the covariance matrix Ω as diffuse, setting them equal to arbitrarily large numbers. We use a smoothing algorithm to obtain Kalman smoothed estimates, conditioned on information from the full sample period (see, e.g., Hamilton (1994)).

5.4 Results

5.4.1 Option Trading and Investor Performance

In this section we compare the raw returns and alphas of option traders and equity investors. Barber and Odean (2000) show that the common stock performance of individual investors is poor due to excessive trading and poor stock selection skills. This raises the question how these investors perform when trading options, which have a more complex payoff structure.

In Table 5.2, Panel A shows that the average option investor loses 1.81% per month in gross terms during the sample period, which is economically large and statistically significant at a 1% level. In contrast, equity investors only lose 0.58%, which is not significant at conventional levels. The second line in Table 5.2 shows that accounting for risk exposures and style tilts explains only part of the poor performance of option traders. Specifically, while the gross alpha of equity traders is close to zero, the risk- and style-adjusted return of option traders is -1.01%. In the third row of Panel A, we adjust returns for time-varying risk and style exposures using the dynamic factor model in equations (5.3.4)-(5.3.6). Although the average dynamic alpha of option traders is 10 basis points higher than their static alpha, they still underperform equity investors by almost one percent per month. Thus, allowing for time variation in risk and investment styles does not eliminate the underperformance of option traders.

The right hand side of Panel A indicates that the performance gap between option traders and stock investors widens when transaction costs are taken into account. In particular, net alphas of equity investors are 2.75% higher than those of option traders. Thus, option investors not only lose due to poor investment decisions but also suffer from higher trading costs. Part of this underperformance can be explained by the brokerage firm's commission fee structure. Trading costs for option contracts consist of a specific amount per contract, whereas trading costs for stocks are based on a fixed amount and a variable part that depends on the value of the transaction. Because option investors tend to trade many small contracts, trading options is more expensive than trading stocks in relative terms.

To shed more light on the poor performance of option traders we split the sample in two subperiods.⁹ The first period, from January 2000 to March 2003, includes the sharp stock market decline after the burst of the tech bubble. In the second subperiod, from April 2003 to March 2006, the market gradually recovers from the crash. Panel B shows that in the first subperiod option investors earn positive gross alphas and actually outperform equity traders. However, Panel C reveals that in the second subperiod option traders underperform other investors by almost two percent in terms of gross alpha and almost four percent in terms of net alpha. These results suggest that option traders miss part of the recovery of the market.

Panel D reports the beta estimates in the static performance evaluation model for both groups of investors. These betas are based on gross returns but loadings based on net returns are similar. The results indicate that the loadings of the portfolios held by option investors on the market and *SMB* factors are significantly lower than those of the equity-only traders. As expected, their exposure to the call option factor is higher. The positive loadings on the *IT* factor imply that, on average, investors' portfolios are tilted towards technology stocks. The subsample analysis shows that investors lower their exposure to the market, momentum, and technology factors but increase their exposure to the size and value factors in the second subperiod. These results are in line with the finding of Kumar (2009a) that individual investors exhibit time-varying style preferences.

Figure 5.1, which traces the evolution of time-varying alphas and betas for option traders, confirms our finding that these investors primarily underperform in the second subperiod. The plot further shows that although factor betas fluctuate heavily, investors often adjust their exposures too late.¹⁰ For instance, in the first subperiod option traders increase their exposure to the call option factor when the market falls, but decrease their call option exposure at the beginning of the second subperiod when the recovery sets in. Interesting, however, is that option traders already lower their exposure to the *IT* factor in 2000 whereas equity traders keep their exposure constant throughout the first subperiod. A possible explanation for the slower response of equity investors to the burst of the tech bubble is that these traders are more prone to the disposition effect than option traders. Brunnermeier and Nagel (2004) show that hedge funds, which are considered to be managed by sophisticated investors, were also riding the bubble but reduced their exposure to the *IT* sector before stock prices collapsed.

Because the poor performance of option traders compared to equity investors is not explained by risk and style tilts, we consider several other potential explanations. First, we account for known determinants of investor performance by relating returns to a number of investor characteristics. We examine the cross-sectional relation between investor performance and characteristics using the Fama and MacBeth (1973) approach. Each month we run a cross-sectional regression of

⁹We do not estimate dynamic models for subperiods because of the small number of time series observations.

¹⁰The likelihood ratio for testing whether factor loadings are constant in the dynamic model is 22.11, rejecting the null hypothesis that betas are constant at a 1% level.

Table 5.2: **Investment Performance of Option Traders and Equity Traders**

This table reports raw returns and alphas of option traders and equity investors. Panel A presents the performance for the full sample period, January 2000 to March 2006, Panel B for the first subperiod from January 2000 to March 2003, and Panel C for the second subperiod from April 2003 to March 2006. The left hand side of each panel shows gross performance and the right hand side displays net performance. *Diff.* is the performance difference between option traders and equity investors. Static alphas are generated by the static model in (5.3.3) and dynamic alphas are the average alphas produced by the dynamic model in (5.3.4)-(5.3.6). Panel D reports estimated factor loadings in the static model. Newey-West *t*-statistics are in parentheses.

	Gross			Net		
	Options	Stocks	Diff.	Options	Stocks	Diff.
<i>Panel A: Performance Full Period: 2000/01 - 2006/03</i>						
Raw return	-1.81 (-2.31)	-0.58 (-0.49)	-1.23 (-1.60)	-4.46 (-6.83)	-1.57 (-1.35)	-2.89 (-3.24)
Static alpha	-1.01 (-1.72)	0.03 (0.05)	-1.04 (-2.00)	-3.72 (-6.30)	-0.97 (-2.08)	-2.75 (-4.65)
Dynamic alpha	-0.93 (-1.75)	0.03 (0.34)	-0.96 (-2.08)	-3.59 (-6.24)	-0.97 (-2.17)	-2.62 (-4.81)
<i>Panel B: Performance Subperiod 1: 2000/01 - 2003/03</i>						
Raw return	-3.03 (-3.02)	-3.70 (-2.52)	0.67 (0.68)	-5.51 (-5.60)	-4.78 (-3.32)	-0.73 (-0.78)
Static alpha	1.32 (1.17)	1.00 (1.23)	0.32 (0.39)	-1.28 (-1.19)	-0.15 (-0.10)	-1.13 (-1.43)
<i>Panel C: Performance Subperiod 2: 2003/04 - 2006/03</i>						
Raw return	-0.48 (-0.58)	2.80 (2.02)	-3.28 (-2.65)	-3.33 (-4.48)	1.90 (1.41)	-5.23 (-4.45)
Static alpha	-1.82 (-2.92)	-0.01 (-0.01)	-1.81 (-2.15)	-4.74 (-7.87)	-0.87 (-1.29)	-3.87 (-4.90)
	Full period		Subperiod 1		Subperiod 2	
	Options	Stocks	Options	Stocks	Options	Stocks
<i>Panel D: Factor Loadings</i>						
R_M	0.94 (5.55)	1.48 (19.39)	1.20 (4.69)	1.46 (10.48)	0.71 (2.52)	1.14 (3.70)
SMB	0.30 (1.21)	0.99 (7.89)	0.22 (0.74)	0.87 (6.87)	0.60 (2.01)	1.05 (3.18)
HML	-0.13 (-0.63)	0.23 (2.31)	-0.10 (-0.39)	0.26 (2.99)	0.10 (0.41)	0.42 (1.62)
MOM	0.16 (1.26)	0.01 (0.08)	0.21 (1.12)	0.15 (2.17)	0.16 (1.16)	-0.36 (-2.65)
ATMC	0.37 (3.24)	0.06 (0.80)	0.66 (2.18)	-0.12 (-0.59)	0.39 (2.58)	0.22 (1.79)
ATMP	0.04 (0.24)	-0.02 (-0.26)	0.08 (0.44)	0.05 (0.53)	-0.15 (-1.08)	-0.21 (-1.15)
IT	0.31 (2.89)	0.41 (6.35)	0.25 (1.44)	0.48 (5.46)	0.17 (0.99)	0.24 (1.28)
Adj. R^2	61.8	89.1	61.3	94.2	62.8	85.3

Figure 5.1: Time-Varying Alpha and Betas for Option Traders

This figure plots the evolution of the Kalman smoothed alpha and betas for option traders over the period January 2000 through March 2006. The estimates are produced by the dynamic factor model in equations (5.3.4)-(5.3.6) and are based on gross returns.

portfolio returns on investor characteristics,

$$R_{it} = \gamma_{0t} + \sum_{l=1}^L \gamma_{lt} Z_{ilt} + \xi_{it}, \quad (5.4.1)$$

where Z_{ilt} is the value of characteristic l for investor i in month t . We then calculate the Fama-MacBeth estimator for the characteristics, which is the time series average of the monthly cross-sectional parameter estimates. We calculate the standard error of the Fama-MacBeth estimator from the time series of these monthly estimates. We perform this analysis for the full sample period as well as the two subperiods defined before to assess the stability of the relation between performance and characteristics in different market conditions.

We include the following characteristics as independent variables: *options* and *stocks&options*, which are two dummy variables equal to one if investor i trades only options or both stocks and options, respectively, in month t ; monthly portfolio *turnover*; *inactive*, a dummy variable equal to one if monthly portfolio turnover is zero; and *woman* and *joint*, which are two gender dummy variables equal to one if the account is held by a woman or jointly by a man and woman. We also include the *value* of the portfolio at the end of the previous month. Because the descriptive statistics in Table 5.1 show that the distributions of turnover and account value are skewed, we trim these characteristics at the 99th percentile and use their logarithmic transformations. We further add the *age*, gross *income*, and *experience* of investors to the regressions.

Table 5.3 reports results from the Fama-MacBeth regressions. Results in the left hand side of the table correspond to regressions with raw returns as dependent variable. The first column shows that even after controlling for investor and portfolio characteristics, option traders continue to earn lower gross returns than those investors who only trade stocks. This is consistent with the observation in Table 5.1 that the demographic, socioeconomic, and portfolio characteristics of option traders and equity investors are very similar. The estimation results for the two subperiods reported in columns two and three confirm the result from Table 5.2 that the underperformance of option traders is concentrated in the second subperiod. The estimates in the next three columns show that also in the presence of control variables the return difference between option and stock investors increases when transaction costs are taken into account. In addition, the coefficient estimates on the control variables have the expected signs. Specifically, we find that trading lowers net performance, in line with the evidence of Barber and Odean (2000), and that portfolios held by women outperform those of men, confirming the results of Barber and Odean (2001). We also find a strong, positive relation between portfolio value and investor performance and between income and performance, which indicates that investors with larger portfolios and higher income outperform smaller investors.

The cross-sectional analysis above does not account for risk and style differences across investor portfolios. The standard approach to identifying the determinants of cross-sectional variation in risk- and style-adjusted returns consists of two steps.

First, for every portfolio a time series regression of returns on risk and style factors is run to obtain its alpha. Second, a cross-sectional regression of these alphas on investor characteristics is estimated. The main drawback of this approach is the errors-in-variables problem that arises because the factor loadings in the first-stage time series regression are estimated with error. Therefore, we introduce a new method that makes it possible to control for risk and style exposures even when the number of return observations for some investors is small. This approach is based on a purged estimator used in the asset pricing literature by Brennan, Chordia, and Subrahmanyam (1998). Each month, we run the cross-sectional regression (5.4.1) of returns on characteristics. We then regress the vector of monthly cross-sectional coefficients for each characteristic, γ_{lt} , on a constant and the risk and style factors F_{kt} ,

$$\gamma_{lt} = \delta_{l0} + \sum_{k=1}^K \delta_{lk} F_{kt} + \omega_{lt}. \quad (5.4.2)$$

In Appendix 5.A we demonstrate that the intercept in this time series regression, δ_{l0} , is an unbiased estimate of the cross-sectional relation between characteristic l and alpha.

The right hand side of Table 5.3 shows that the conclusion that option trading has a detrimental impact on investor performance also holds when alpha is used as a dependent variable. The bottom line is that controlling for differences in risk and style exposures and differences in the characteristics of investors and their portfolios does not explain the underperformance of option traders compared to equity investors.

5.4.2 Investor Sentiment and Market Timing

The finding that the underperformance of option investors is concentrated in the second subperiod suggests that it is related to poor market timing. Market timing skills are especially important in the option market because options are effective instruments for betting on market moves due to the leverage they provide and the positive skewness of their payoffs. Since we find that option investors perform poorly when markets move upward, we conjecture that after the stock market decline in 2001 and 2002, investors became bearish about market prospects and expected a further fall. Because the brokerage firm restricts the ability of its clients to sell stocks short to the 50 largest and most liquid Dutch stocks and only allows intraday short selling, most investors use options to speculate on a market decrease. Therefore, we expect that at the end of 2002 investors took bearish positions in options.

Initial evidence supporting this hypothesis is provided in Figure 5.2, which plots the return on the MSCI Netherlands equity index, the gross return difference between equity investors and option traders, and the ratio of short-term (three months to expiration or less) option positions taken in anticipation of a market increase and positions taken in anticipation of a market decrease. We calculate

Table 5.3: Investor Performance and Characteristics

This table reports Fama and MacBeth (1973) estimates for the full sample, January 2000 to March 2006, and for two subperiods, January 2000 to March 2003 and April 2003 to March 2006. The left side of the table uses raw returns as dependent variables and the right side refers to alphas. The independent variables are investor characteristics. *Options* and *Stocks&options* are dummy variables equal to one if an investor trades only options or both stocks and options. *Turnover* is the log of monthly portfolio turnover. *Inactive* is a dummy equal to one if an investor does not trade. *Single woman* and *Joint* are dummies equal to one if the account is held by a woman or jointly by a man and woman. *Portfolio value* is the log of portfolio value and is lagged by one month. *Age*, *Income*, and *Experience* are as defined in Table 5.1. *t*-Statistics are in parentheses.

	Raw Return						Alpha					
	Gross			Net			Gross			Net		
	Full	Sub 1	Sub 2	Full	Sub 1	Sub 2	Full	Sub 1	Sub 2	Full	Sub 1	Sub 2
Options	-1.34 (-1.88)	0.40 (0.37)	-3.23 (-3.95)	-3.38 (-4.73)	-1.39 (-1.28)	-5.55 (-7.10)	-1.11 (-2.08)	0.39 (0.41)	-1.96 (-2.77)	-3.21 (-6.09)	-1.54 (-1.58)	-4.32 (-6.68)
Stocks&options	-0.91 (-3.96)	-1.10 (-3.09)	-0.70 (-2.47)	-1.64 (-7.04)	-1.77 (-5.01)	-1.50 (-4.97)	-0.79 (-3.61)	-0.69 (-1.69)	-0.64 (-2.25)	-1.53 (-6.91)	-1.39 (-3.45)	-1.42 (-4.92)
Turnover	0.67 (4.54)	0.84 (3.34)	0.47 (3.57)	-0.08 (-0.63)	0.08 (0.36)	-0.27 (-2.20)	0.79 (5.72)	1.38 (5.05)	0.29 (2.10)	0.04 (0.28)	0.54 (1.16)	-0.42 (-2.96)
Inactive	-1.16 (-2.83)	-0.92 (-1.57)	-1.42 (-2.47)	0.79 (2.03)	1.05 (1.87)	0.52 (0.95)	-1.46 (-4.26)	-2.07 (-3.42)	-0.27 (-0.57)	0.52 (1.60)	-0.02 (-0.03)	1.64 (3.60)
Single woman	0.38 (3.47)	0.67 (3.96)	0.07 (0.60)	0.37 (3.35)	0.66 (3.96)	0.06 (0.45)	0.33 (3.52)	0.36 (2.12)	0.19 (1.52)	0.31 (3.32)	0.34 (2.06)	0.16 (1.28)
Joint	0.18 (3.44)	0.19 (2.12)	0.16 (3.53)	0.23 (4.44)	0.25 (2.83)	0.20 (4.39)	0.13 (2.29)	-0.02 (-0.21)	0.14 (3.89)	0.18 (3.75)	0.04 (0.42)	0.17 (2.76)
Portfolio value	0.28 (2.52)	0.38 (2.03)	0.17 (1.96)	0.55 (5.21)	0.66 (3.48)	0.44 (5.35)	0.22 (2.01)	0.02 (0.07)	0.23 (2.45)	0.50 (5.07)	0.31 (1.48)	0.41 (3.97)
Age/10	-0.01 (-0.28)	0.01 (0.22)	-0.03 (-1.03)	-0.02 (-0.59)	0.01 (0.23)	-0.05 (-1.78)	-0.05 (-2.32)	-0.05 (-1.39)	-0.05 (-2.04)	-0.07 (-2.78)	-0.06 (-1.14)	-0.07 (-2.19)
Income	0.47 (5.16)	0.73 (5.42)	0.19 (1.80)	0.50 (5.55)	0.77 (5.81)	0.21 (2.03)	0.43 (3.54)	0.84 (4.09)	0.35 (5.25)	0.46 (4.87)	0.89 (5.67)	0.37 (3.50)
Experience	0.07 (1.09)	0.12 (0.95)	0.02 (0.42)	0.07 (1.09)	0.11 (0.92)	0.02 (0.43)	0.11 (1.04)	0.38 (1.23)	0.02 (0.79)	0.05 (0.84)	0.16 (1.11)	0.03 (0.20)

this call/put ratio as the value of call options bought by the investors divided by the value of their put option purchases.¹¹ An increase (decrease) in the call/put ratio indicates that option investors become more optimistic (pessimistic) about stock market prospects. Figure 5.2 indicates that the ratio is below one for most months in 2003 and 2004, which is exactly the period in which stock markets started to recover. This supports our hypothesis that after the stock market crash in 2001 and 2002, option traders speculated on a further decline of the market. As a result, they missed part of the recovery of the market and consequently, common stock investors outperformed option traders during this period.

The importance of sentiment in option markets has recently been shown by Han (2008), who documents that investor sentiment about the stock market affects option prices. Figure 5.3 relates the call/put ratio to the level of the MSCI Netherlands index and to two other sentiment indicators, which are the consumer confidence index obtained from the Dutch Central Bureau of Statistics and the AEX volatility index (VIX) provided by Euronext. The VIX measures the market's expectation of short-term volatility implied by AEX index option prices. Since volatility often signals financial turmoil, the VIX is commonly interpreted as a measure of fear in the market.¹² We plot the inverse of the VIX so that high values correspond to optimism among investors. Figure 5.3 shows a strong relation between the MSCI index on the one hand and the three sentiment measures on the other hand. Both the index and the sentiment indicators reach their lowest value in the beginning of 2003, which coincides with the start of the second subperiod. The plots also reveal that after the fall of the market it takes considerable time before confidence of consumers and investors has been restored to normal levels.

Table 5.4 provides a more formal analysis of the relation between sentiment, market timing, and investor performance. The pairwise correlations between the three different sentiment measures reported in Panel A confirm their strong relationships depicted in Figure 5.3. In Panel B we present results for time series regressions of the gross return difference between option traders and equity investors on the market return, the call/put ratio, and interaction terms between these two variables. The first column shows that the market return alone explains more than 37% of the performance difference between these two groups of investors. The negative coefficient implies that option traders lose relative to stock investors when markets increase, which lends further support to the hypothesis that bad market timing is responsible for the underperformance of option traders in the second subperiod. The second column indicates that adding the call/put ratio only leads to a small increase in the adjusted R^2 . In contrast, column 3 shows that adding an interaction term between the market return and the call/put ratio leads to a sharp increase in explanatory power. In fact, this two-factor model explains more than 60% of the underperformance of option traders. The positive

¹¹We also considered an extension of this ratio, defined as the sum of the value of call options bought and put options sold divided by the sum of the value of put options bought and call options sold. Using this ratio leads to similar results.

¹²The construction of the AEX volatility index follows the methodology of the CBOE for constructing the VIX in the U.S., which is based on S&P 500 index option prices.

Figure 5.2: Call/Put Ratio and Return Difference between Equity Traders and Option Traders

This figure plots the evolution through time of the monthly call/put ratio, the return on the MSCI Netherlands equity index, and the gross return difference between equity and option traders. We calculate the call/put ratio as the value of short-term (three months to expiration or less) call options bought by the investors divided by the value of their put option purchases. The sample period is January 2000 to March 2006.

coefficient on this interaction term means that for a given market return, an increase in the call/put ratio is related to an increase in the relative performance of option traders. This explains the poor performance of option investors in the second subperiod, when the call/put ratio of their trades was low.

In columns four and five we include call/put ratios constructed using index options and stock options, respectively, to examine whether the underperformance of option traders is driven by market sentiment (index options) or stock-specific sentiment (stock options). When including these variables separately, they are both significant at a 1% level. In column six we include both ratios simultane-

Figure 5.3: MSCI Netherlands Index and Sentiment Measures

This figure shows the evolution of the MSCI Netherlands Equity Index, scaled to 1 in January 2000, and of three sentiment measures. We define the call/put ratio as the value of short-term (three months to expiration or less) call options bought by the investors divided by the value of their put option purchases. The AEX VIX index captures the implied volatility embedded in prices of options on the Dutch AEX stock market index. We plot the inverse of the VIX so that low values correspond to fear among investors. The Consumer Confidence Index is constructed by the Dutch Central Bureau of Statistics. The sample period is from January 2000 to March 2006.

ously in the regression and find that they are still significant. This suggests that market sentiment and stock-specific sentiment are both important in explaining the underperformance of option traders.

The conclusion that option traders overreact to past stock market movements is in line with results of Lakonishok, Lee, Pearson, and Poteshman (2007). These authors find that discount clients decreased their open interest in purchased puts at the height of the Internet bubble but sharply cut their bets that prices would increase after the burst of the bubble. In contrast, clients of full-service brokers and firm proprietary traders did not use the option market to speculate on rising prices during the boom and a further fall in prices after its collapse.

Table 5.4: **Sentiment and Market Timing**

Panel A of this table presents descriptive statistics for three sentiment indicators. We define the *Call/Put (C/P) ratio* as the value of all call options bought by the investors divided by the value of their put option purchases. The *AEX VIX index* captures the implied volatility embedded in prices of options on the Dutch AEX stock market index. The *CC index* is the consumer confidence index constructed by the Dutch Central Bureau of Statistics (CBS). Panel B presents results for time series regressions of the gross return difference between option traders and equity traders on the market return, the call/put ratio, and various interaction terms between the market return and the call/put ratio, constructed using all options (*C/P ratio*), index options only (*C/P index*), or stock options only (*C/P stocks*). *t*-Statistics based on Newey-West heteroskedasticity and autocorrelation robust standard errors are in parentheses.

	Mean	Std. Dev.	AR(1)	Pairwise Correlations		
<i>Panel A: Descriptive Statistics on Sentiment Indicators</i>						
C/P ratio	1.12	0.49	0.79	1.00		
AEX VIX index	24.63	10.77	0.86	−0.60	1.00	
CC index	−13.45	19.74	0.97	0.43	−0.52	1.00
Gross Return Difference Option and Equity Traders						
<i>Panel B: Option Trader Performance and Market Timing</i>						
Intercept	−1.91 (−3.08)	−4.48 (−2.95)	−1.99 (−4.16)	−1.82 (−3.67)	−1.97 (−4.00)	−1.90 (−4.01)
Market	−0.77 (−6.57)	−0.78 (−6.74)	−2.14 (−10.00)	−3.29 (−8.20)	−2.35 (−9.10)	−3.15 (−8.18)
C/P ratio		2.31 (1.85)				
Market×C/P ratio			1.46 (7.06)			
Market×C/P index				3.14 (6.45)		1.79 (2.70)
Market×C/P stocks					0.66 (6.55)	0.40 (2.87)
Adj. R^2	37.1	38.3	61.8	59.0	59.5	62.8

To further investigate the relation between sentiment and market returns, we estimate vector autoregressive (VAR) models with the call/put ratio and the market return as dependent variables. The VAR model makes it possible to examine the dynamic interrelationships between sentiment and market returns. The methodology is similar to that used by Brown and Cliff (2004) and Schmitz, Glaser, and Weber (2009), who find strong evidence that sentiment is driven by past stock market returns but little evidence that sentiment predicts returns.

The general VAR model is given by

$$Y_t = \phi_0 + \sum_{j=1}^P \phi_j Y_{t-j} + \eta_t, \quad (5.4.3)$$

where Y_t is a vector that contains the call/put ratio and the market return. We use the Schwarz criterion and the Akaike information criterion to determine the appropriate number of lags P . These criteria prefer a specification with two lags of the call/put ratio and the market return.

Table 5.5 presents the estimation results from the VAR model for the full sample period and the two subperiods. The first column shows that the first-order lags of the call/put ratio and the market return have a significantly positive impact on the current level of the call/put ratio. This is consistent with the hypothesis that option investors extrapolate recent market returns. In contrast, Schmitz, Glaser, and Weber (2009) derive a sentiment measure from warrants and find that past returns have a negative impact on sentiment, while sentiment has a positive impact on future returns. However, these effects are short-lived (2 days), so sentiment measures do not seem to be useful to predict returns for longer horizons. Table 5.5 also reports F -statistics and p -values for Granger causality tests that the coefficients on the lags of the independent variables other than the lags of the dependent variable are jointly equal to zero. The results show that market returns Granger-cause the call/put ratio at the 1% level. In contrast, column two indicates that lags of the call/put ratio do not predict the return on the market. Estimation results for the two subperiods are similar and confirm the conclusion that sentiment among option investors, as measured by the call/put ratio, is driven by past market returns while the call/put ratio itself does not predict future market returns.

5.4.3 Hedging and Gambling Motivations for Option Trading

The large losses on option investments bring up the question why individual investors trade options. In this section we consider two alternative motivations for option trading. First, it is possible that investors use options for hedging the downward risk of their stock portfolio. Another possibility is that investors use options for gambling. Kumar (2009b) links the preference of individual investors for lottery stocks to their socioeconomic characteristics. He defines lottery stocks as stocks with a low price, high idiosyncratic skewness, and high idiosyncratic volatility and shows that state lotteries and lottery stocks attract a similar socioeconomic clientele, which consists of poor, young, single men. However, options seem to be even more attractive for gambling purposes than stocks because of the leverage they provide and their skewed payoffs.

We examine this hypothesis by linking both option trading and equity trading to a number of investor and portfolio characteristics. Specifically, we regress different measures of trading on a gender dummy variable, lagged portfolio value, age, income, experience, and past portfolio returns. We also include the number of trades in the other type of instrument (stocks or options) as a regressor to control for an investor's general propensity to trade. Furthermore, this allows us to investigate whether option investors also trade more stocks than equity-only investors.

Table 5.5: VAR Model of Sentiment and Market Returns

This table presents estimation results for vector autoregressive models with two lags. The dependent variables are the *Call/Put (C/P) ratio* and the *Market* return. We estimate the VAR models for the full period, January 2000 to March 2006, and for two subperiods, January 2000 to March 2003 (subperiod 1) and April 2003 to March 2006 (subperiod 2). *t*-Statistics based on Newey-West heteroskedasticity and autocorrelation robust standard errors are in parentheses. *Granger F* is the *F*-statistic of a test that the coefficients on the lags of the independent variables other than the lags of the dependent variable are jointly equal to zero.

	Full Period		Subperiod 1		Subperiod 2	
	C/P Ratio	Market	C/P Ratio	Market	C/P Ratio	Market
Intercept	0.26 (3.71)	0.89 (0.55)	0.29 (2.28)	-3.58 (-1.29)	0.18 (1.41)	2.02 (0.80)
C/P_{t-1}	0.63 (6.13)	1.10 (0.47)	0.67 (4.60)	2.98 (0.95)	0.45 (2.64)	-0.26 (-0.08)
C/P_{t-2}	0.11 (1.15)	-2.03 (-0.91)	0.05 (0.40)	-1.92 (-0.64)	0.35 (2.10)	0.68 (0.21)
$Market_{t-1}$	0.02 (3.09)	0.06 (0.53)	0.02 (1.96)	-0.06 (-0.36)	0.02 (2.48)	-0.21 (-1.25)
$Market_{t-2}$	-0.01 (-0.84)	0.10 (0.81)	-0.01 (-0.62)	-0.02 (-0.09)	0.00 (0.35)	-0.14 (-0.80)
Granger <i>F</i>	5.08	0.53	2.28	0.48	3.09	0.03
<i>p</i> -Value	(0.01)	(0.59)	(0.12)	(0.62)	(0.06)	(0.97)
Adj. R^2	70.3	-2.4	70.5	-9.1	54.7	-5.2

The first two columns in Table 5.6 report estimates for a pooled probit model that measures the influence of the investor and portfolio characteristics on the decision to trade options (column one) and stocks (column two) in a given month. Standard errors are clustered by investor and by time period, as suggested by Petersen (2009), to account for both serial correlation and cross-sectional correlation. The results indicate that past portfolio returns have a significantly negative impact on the decision to trade options. In contrast, past returns have a positive but insignificant effect on the decision to trade stocks. The finding that lower past returns increase trading activity among option investors but decrease the intensity of trading among equity investors is consistent with a disposition effect being present among equity traders but absent among option traders. Another potential explanation for this finding is that option traders are loss-averse and take more risk following losses and less risk following gains. This is consistent with evidence documented by Coval and Shumway (2005), who show that professional futures traders are highly loss-averse.

Table 5.6 also shows that females have a lower propensity to trade both options and stocks, which generalizes the conclusion of Barber and Odean (2001) that men trade more than women to the option market. Portfolio value has a positive impact on the decision to trade, since large portfolios require more trades. We further find that older investors are more likely to trade stocks and options. A

possible explanation for this result is that for older investors trading is a leisure activity.¹³ Experience and income lower the probability of trading stocks and options, consistent with the view that less experienced investors and those with lower income are more inclined to gamble.

The probit regressions consider the decision to trade but not the intensity of trading. Therefore, in columns three and four we report Fama and MacBeth (1973) regression estimates in which the dependent variable is the logarithm of the number of option trades (column three) and equity trades (column four). Since the dependent variables exhibit serial correlation, we correct the Fama-MacBeth standard errors for serial correlation using the Newey-West procedure, as suggested by Cochrane (2005). In addition, we apply the Heckman (1976) two-stage procedure to account for self-selection in the trading decision. In the first stage, we use the estimates from the probit regressions to calculate the inverse Mills ratio. We then include the inverse Mills ratio as an additional regressor in the Fama-MacBeth regressions to obtain consistent estimates.

The Fama-MacBeth estimates are qualitatively similar to the results from the probit model. In particular, higher portfolio returns in the previous month decrease option trading activity but increase the intensity of equity trading. Single men, investors with large accounts, older investors, and those with low income and little experience tend to trade most. Furthermore, the results confirm the hypothesis that investors who trade more options (stocks) also trade more stocks (options). Although we correct the Fama-MacBeth standard errors for serial correlation, as a robustness check we also report estimation results for pooled panel regressions in the last two columns with standard errors clustered by investor and by time period. The results from the panel model are very similar to the Fama-MacBeth estimates and confirm the conclusion that investor and portfolio characteristics have the same influence on option trading as on equity trading, except for the impact of past portfolio returns. The finding that single men with low income and little experience trade most, suggests that gambling and entertainment play an important role in explaining excessive trading by individual investors in both option and stock markets.

To dig further into the motivations for trading options, in Table 5.7 we classify all opening option trades into purchased call, written call, purchased put, and written put positions. Furthermore, a split is made between AEX index options and stock options. Panel A reports the percentage of option volume in a given category relative to the total of calls and puts traded. When looking at trading volume in all options, purchased call options dominate, followed by written calls, written puts, and purchased put positions. The result that call positions are more prevalent than put positions is consistent with results documented by Lakonishok, Lee, Pearson, and Poteshman (2007). However, differentiating between index options and stock options shows that purchased put positions dominate for index options but are least prevalent for stock options. Similarly, Schmitz, Glaser, and Weber (2009) find that for index warrants, trading volume of calls equals volume of puts but that for stock warrants, volume of calls is much larger than volume

¹³ Another reason could be that their risk aversion is lower because they tend to have more wealth.

Table 5.6: **Trading Behavior and Investor Characteristics**

This table relates the trading behavior of individual investors to investor characteristics. Results in the columns labeled “Probit” are estimates for a pooled probit regression. Standard errors clustered by investor and month to account for potential serial correlation and cross-sectional correlation in residuals are in parentheses. Results in the columns labeled “Fama-MacBeth” refer to Fama and MacBeth (1973) coefficient estimates with Newey-West heteroskedasticity and autocorrelation robust standard errors in parentheses. Results in the columns labeled “Panel” are pooled OLS panel regression estimates with standard errors clustered by investor and month in parentheses. The dependent variable in the probit model is a dummy variable equal to one if an investor trades options (column 1) or stocks (column 2) in a given month. The dependent variable in the Fama-MacBeth and panel regressions is the logarithm of the number of option trades per month per account (columns 3 and 5) or the logarithm of the number of equity trades (columns 4 and 6). The inverse Mills ratio is based on the estimates of the pooled probit regressions and is included to account for self-selection in trading decisions. The R^2 for the probit regressions is the pseudo R^2 .

	Probit		Fama-MacBeth		Panel	
	Options	Stocks	Options	Stocks	Options	Stocks
Intercept	-1.51 (-10.63)	-1.09 (-5.19)	-2.40 (-6.14)	0.20 (0.91)	-3.12 (-4.56)	0.02 (0.06)
Net return _{t-1}	-0.33 (-4.80)	0.06 (1.14)	-0.81 (-10.06)	0.02 (0.52)	-0.76 (-7.15)	0.04 (0.68)
Single woman	-0.20 (-9.15)	-0.23 (-7.56)	-0.33 (-8.79)	-0.21 (-11.26)	-0.40 (-6.34)	-0.25 (-8.02)
Joint	-0.13 (-10.00)	-0.07 (-3.61)	-0.22 (-11.11)	-0.06 (-8.34)	-0.26 (-6.99)	-0.07 (-5.39)
Portfolio value _{t-1}	0.16 (38.35)	0.23 (20.58)	0.35 (14.56)	0.25 (14.64)	0.38 (8.86)	0.27 (10.86)
Age/10	0.04 (7.57)	0.00 (0.22)	0.09 (12.02)	0.02 (6.84)	0.10 (8.48)	0.02 (4.98)
Income	-0.12 (-6.49)	-0.17 (-7.09)	-0.29 (-16.23)	-0.25 (-27.27)	-0.30 (-7.86)	-0.26 (-11.42)
Experience/12	-0.08 (-14.17)	-0.01 (-5.72)	-0.25 (-7.12)	-0.17 (-4.17)	-0.17 (-7.42)	-0.12 (-9.44)
Inverse Mills			2.13 (10.99)	0.79 (8.87)	2.46 (7.12)	0.97 (6.63)
Equity trades	0.22 (30.35)		0.48 (16.12)		0.53 (9.13)	
Option trades		0.23 (9.03)		0.18 (11.72)		0.20 (15.13)
R^2	6.6	9.2	6.3	6.9	4.3	5.6

of puts. The finding that purchased puts are the least frequently traded category of stock options suggests that most investors do not use options for hedging their underlying stock portfolio. Panel B shows the percentage of the value of option trades in a given category relative to the total value of calls and puts traded. These results are qualitatively similar to those reported in Panel A. In particular, most investments in index options are purchased put positions while for stock options the highest amount of money is invested in long call positions.

In Panel C we display the moneyness of the options in each category, defined as the strike price divided by the price of the underlying asset. The results indicate that investors have a strong preference for out-of-the-money options. This holds for both calls and puts and for index options and stock options. The tendency to take out-of-the-money option positions is consistent with a gambling motive for option trading, since these options are cheap and offer a small probability of a large gain, similar to a lottery ticket.

Panel D reports the percentage of stock options for which the investor holds the underlying stock. These results provide direct evidence that most investors do not use options for hedging. In fact, more than 70% of all purchased puts are naked positions. The option category with the highest percentage of options for which the investor has a long position in the underlying stock is that of written calls, which suggests that some investors use a strategy of covered call writing. Although written calls can be used for hedging, the protection they provide against a price decline of the underlying asset is limited to the option premium. Given the restrictions imposed by the broker on short selling stocks, it is unlikely that investors use purchased call and written put positions for hedging short stock positions. Panel D indeed shows that most positions taken in these categories are also naked positions. Our conclusion that most individual investors primarily use options to speculate on directional price movements reinforces results documented by Lakonishok, Lee, Pearson, and Poteshman (2007) for the U.S. market.¹⁴

We use survey data to complement the information derived from the trades and portfolio positions of the investors. In September 2005, a questionnaire was sent to all clients of the brokerage firm. In total 4,516 clients responded, which we split into 2,323 option traders and 2,193 equity-only traders. Panel A in Table 5.8 reports the responses to statements related to investor experience and to the importance of the portfolio held at the online broker. An important difference between option traders and equity investors is that only a small proportion of option traders consider themselves novice investors while more than half of all equity traders think they are novice investors. However, the average option investor indicates that she has been trading for just four years, which is only slightly longer than the investment experience of equity traders. This suggests that option traders are overconfident about their investment experience. Panel A further shows that both option and equity traders indicate that they invest more than half of their portfolio at the broker. This implies that the losses investors incur on accounts held at the broker can have a serious impact on their financial wealth.

¹⁴A preliminary analysis also confirms their conclusion that only a small proportion of option trading involves straddles, which can be used to trade volatility.

Table 5.7: **Descriptive Statistics on Option Positions**

In this table we classify all 2,497,378 opening option trades into purchased call, written call, purchased put, and written put positions. Furthermore, a split is made between index options (1,148,148 trades) and stock options (1,349,230 trades). Panel A reports the percentage of option volume in a given category relative to the total of calls and puts traded. Panel B shows the percentage of the value of option trades in a given category relative to the total value of calls and puts traded. In Panel C we display the moneyness of the options in each category, defined as the strike price K divided by the price of the underlying asset S . Panel D reports the percentage of stock options for which the investor holds the underlying stock.

	Call			Put		
	Purchased	Written	Total	Purchased	Written	Total
<i>Panel A: Option Volume (%)</i>						
All	40.4	21.3	61.7	17.9	20.3	38.3
Index	30.3	17.4	47.6	33.1	19.2	52.4
Stock	45.2	23.2	68.4	10.8	20.8	31.6
<i>Panel B: Option Value (%)</i>						
All	34.4	16.9	51.3	28.1	20.6	48.7
Index	26.3	16.2	42.5	37.9	19.5	57.5
Stock	49.6	18.3	67.8	9.5	22.7	32.2
<i>Panel C: Option Moneyness (K/S)</i>						
All	1.02	1.03	1.02	0.97	0.94	0.96
Index	1.03	1.03	1.03	0.98	0.95	0.97
Stock	1.01	1.03	1.02	0.96	0.94	0.95
<i>Panel D: Underlying Stock Holdings (%)</i>						
Yes	14.5	34.3	21.1	27.3	27.1	27.2
No	85.5	65.7	78.9	72.7	72.9	72.8

Panel B shows the responses to statements that measure the investment attitude of the clients. More than half of all option traders indicate that investing is just a hobby for them. Consistent with this result, over 80% of the option investors and 65% of the equity traders state that they only invest their spare money. Moreover, whereas 17.2% of the option traders agree to the statement that short-term speculation is their main investment objective, only 9.4% of the equity investors agree. The overall picture that emerges from the results reported in Panel B is that for both groups of investors entertainment and sensation-seeking are important reasons for investing. However, option traders seem to be affected most by these factors.

Panel C presents the responses to three statements on option trading. As expected, only a small proportion of option investors and a majority of equity-only traders consider option trading too risky. A striking result is that only 10% of those who do invest in options indicate that they use the option Greeks when

trading options. This finding suggests that a lack of knowledge about the risk and return characteristics and the use of options might also contribute to the poor performance of option investments. Finally, Panel C reveals that only 20% of the option traders indicate that they use options primarily for hedging their risks, which confirms the conclusion from Table 5.7 that little option volume can be attributed to hedging.

Table 5.8: **Survey on Investment Attitude and Option Trading**

This table shows the responses of 4,516 clients of the brokerage firm to statements on investor experience (Panel A), investment attitude (Panel B), and option trading (Panel C). We split the respondents into 2,323 option traders and 2,193 equity traders. The questionnaire was sent to all clients of the brokerage firm in September 2005.

	Option Traders	Equity Traders
<i>Panel A: Investor Experience (% agree or amount x)</i>		
"I consider myself a novice investor."	18.2	50.7
"I already invest for x years."	4.0	3.5
"I invest $x\%$ of my portfolio at another bank or broker."	48.8	46.0
<i>Panel B: Investment Attitude (% agree)</i>		
"Investing is just a hobby for me."	55.7	43.6
"I only invest my spare money."	80.9	65.3
"My main investment objective is short-term speculation."	17.2	9.4
<i>Panel C: Option Trading (% agree)</i>		
"Option trading is too risky."	11.6	67.8
"When trading options I use the option Greeks."	11.2	—
"I primarily use options to hedge my risks."	19.6	—

5.4.4 Option Trading Skills

Although the previous sections have shown that on average option traders perform poorly, it is possible that some sophisticated investors do possess trading skills and are able to exploit inefficiencies in the option market. Coval, Hirshleifer, and Shumway (2005) present strong evidence of performance persistence among a small group of common stock investors. They note that although it is unlikely that individual investors are better informed than professional fund managers, they are better able to exploit superior information for two reasons. First, individual investors usually trade smaller positions, so the price impact of their trades is limited. Second, individual investors face fewer asset allocation constraints, since they are not required to hold a diversified portfolio or track a specified benchmark. In this section we perform persistence tests to investigate whether their findings can be extended to option traders. Subsequently, we examine the characteristics

and trading behavior of successful and unsuccessful option investors.

We first sort option traders into decile portfolios based on their performance during a ranking period. We then calculate returns for each of these deciles over an evaluation period. Repeating the ranking procedure using nonoverlapping intervals produces a time series of post-ranking returns for each decile. We keep investors who drop out of the sample during the evaluation period in the decile portfolios until they disappear, after which we adjust the portfolio weights. We test whether past winners continue to outperform past losers by performing a *t*-test on the return difference between decile 1 (past winners) and decile 10 (past losers). We also calculate Spearman rank correlation coefficients between the formation period ranking and the evaluation period ranking. The null hypothesis of the Spearman test is that there is no relation between formation and evaluation period ranking, i.e., no performance persistence.

To examine whether consistency in returns, if any, is due to persistence in transaction costs, we perform the analysis using both gross and net returns. If we only find evidence of performance persistence when we sort on net returns, we conclude that it is related to costs. We consider three-, six-, and 12-month ranking and evaluation periods. This choice is based on mutual fund studies showing that performance persistence is usually short-term (e.g., Busse and Irvine (2006)). Using longer periods also means that we can include fewer investors in the analysis, because we require an investor to be present in the sample during the complete ranking period and at least one month in the evaluation period. On the other hand, results for periods shorter than three months are likely to be dominated by noise and luck.

Table 5.9 presents post-formation returns and alphas for portfolios of option investors sorted on past one-year return. The columns labeled “gross” (“net”) refer to deciles formed and evaluated on the basis of gross (net) performance. The results indicate that on average, option traders in the top decile continue to outperform those in the bottom decile in the year subsequent to the formation year by more than 3.5% per month in terms of gross return and almost 5% in terms of net return. The consistency in gross returns indicates that persistence in costs explains only part of total performance persistence. A substantial part of the performance differential is driven by the extremely poor performance of the option investors in decile 10. Gross and net returns for these investors are 2% and 2.5% lower, respectively, than for those in decile 9. Spearman rank correlations (not reported to save space) are 0.93 and significant at the 1% level, which indicates a strong relation between formation and evaluation period ranking.¹⁵

The difference in performance between past winners and losers also shows up when we consider risk- and style-adjusted returns. Investors in decile 1 earn gross and net alphas in the evaluation period that are almost 4% and 5% higher per month, respectively, than those earned by investors in decile 10. These differences are significant at a 1% level. Spearman rank correlation coefficients exceed 0.9 and are also significant at the 1% level, providing strong evidence of performance

¹⁵Results for three- and six-month ranking and evaluation periods (omitted to save space) show that performance persistence is still significant on six-month horizons but not for shorter periods.

persistence among option traders. Although gross and net alphas for option investors in the top deciles are positive but insignificant, alphas are significantly negative for those in the bottom two deciles. Table 5.9 also reports the exposures of the returns on the decile portfolios to the factors in the performance attribution model. Loser deciles tend to have significantly higher loadings on the *SMB* and *IT* factors than winner deciles, indicating strong tilts in their portfolios towards small tech stocks.

Since the performance difference between winners and losers cannot be fully explained by trading costs and risk and style tilts, we investigate whether it can be linked to heterogeneity in the characteristics of investors and their portfolios. Table 5.10 presents the characteristics of the investors in the decile portfolios at the end of the ranking period. We find that the median account value increases with ranking, i.e., the top decile consists of investors who hold the largest accounts. Furthermore, those in the bottom deciles have much higher turnover than successful option traders. Table 5.10 also shows that in the bottom deciles a higher proportion of accounts is held by men with low income and little experience.

5.5 Conclusion

This chapter shows that option trading has a detrimental impact on the performance of individual investors. The losses investors incur on their option investments are much larger than those from equity trading. Risk and style tilts and differences in demographic and socioeconomic characteristics do not explain the poor performance of option traders relative to equity investors. Instead, we attribute the poor performance of option traders to bad market timing due to overreaction to past stock market movements. In particular, after the collapse of the Internet bubble, option traders speculated on a further market decrease when markets actually started to recover. High trading costs also contribute to the losses suffered by option investors.

We also show that various demographic, socioeconomic, and portfolio characteristics that have been linked to gambling by Kumar (2009b), have a similar influence on the trading intensity of option investors and equity investors. Specifically, we find that single men with low income and little investment experience are most likely to engage in both option trading and equity trading. However, an important difference between option traders and equity investors is that the trading activity of the former group increases after past losses while past performance has a positive, but insignificant effect on the trading volume of equity investors. By linking option trades to the common stock holdings of individual investors, we rule out hedging as an important motivation for trading options. Instead, investors tend to take naked, out-of-the-money option positions, suggesting a gambling motive for trading options. Options are particularly attractive for gambling because of the leverage they provide and their skewed payoffs, thereby resembling lottery tickets. Responses of investors to statements on investment attitude confirm that entertainment and sensation-seeking are important reasons for trading options.

Table 5.9: Performance Persistence of Option Traders

At the end of every year we sort option investors into equal-weighted decile portfolios based on returns earned over the year. Each portfolio is held for one year and subsequently rebalanced. This table shows the average monthly raw returns, alphas, and factor loadings for each decile portfolio in the post-formation period. Decile 1 contains the 10% of investors with the highest return during the ranking period and decile 10 includes the worst 10% performers in the ranking period. Columns labeled “Gross” (“Net”) refer to deciles formed and evaluated based on gross (net) returns. *t*-Statistics based on Newey–West heteroskedasticity and autocorrelation robust standard errors are in parentheses.

Decile	Gross		Net		Factor Loadings							Adj. R^2
	Return	Alpha	Return	Alpha	R_M	SMB	HML	MOM	ATMC	ATMP	IT	
1	0.07 (0.08)	0.22 (0.73)	-0.05 (-0.07)	0.03 (0.13)	1.03 (11.35)	0.16 (1.07)	0.13 (1.27)	0.01 (0.22)	-0.11 (-0.15)	-1.85 (-2.25)	-0.06 (-1.31)	84.8
2	0.39 (0.51)	0.49 (2.34)	-0.09 (-0.12)	0.23 (1.18)	0.81 (13.52)	0.17 (2.53)	0.20 (4.59)	-0.10 (-2.78)	-1.13 (-2.60)	0.55 (0.97)	0.05 (1.01)	91.5
3	-0.07 (-0.10)	0.16 (0.90)	0.02 (0.03)	0.08 (0.37)	0.79 (23.44)	0.13 (4.29)	-0.05 (-1.15)	-0.02 (-0.99)	-0.04 (-0.17)	-0.43 (-1.44)	0.06 (2.25)	94.4
4	0.28 (0.35)	0.44 (1.82)	-0.12 (-0.16)	0.03 (0.14)	0.97 (25.32)	0.38 (9.81)	0.14 (3.39)	0.05 (1.49)	0.80 (2.68)	0.32 (0.05)	0.21 (5.46)	95.2
5	-0.08 (-0.09)	0.09 (0.31)	-0.44 (-0.45)	-0.30 (-0.97)	1.19 (17.56)	0.47 (4.20)	0.10 (1.18)	0.03 (0.52)	-1.29 (-3.22)	0.82 (0.13)	0.19 (3.75)	94.0
6	-0.31 (-0.32)	0.02 (0.08)	-0.48 (-0.50)	-0.23 (-0.79)	1.16 (18.32)	0.57 (3.72)	-0.01 (-0.14)	0.11 (1.53)	0.57 (0.89)	-1.21 (-1.82)	0.32 (5.18)	87.6
7	-0.42 (-0.44)	-0.13 (-0.35)	-0.87 (-0.84)	-0.77 (-1.44)	1.25 (19.22)	0.90 (8.95)	0.15 (1.50)	0.25 (3.99)	-0.76 (-1.40)	1.71 (2.23)	0.39 (6.61)	87.2
8	-0.39 (-0.32)	-0.17 (-0.34)	-1.05 (-0.91)	-0.68 (-1.14)	1.51 (18.83)	1.17 (12.35)	0.08 (0.70)	0.27 (3.99)	0.37 (0.59)	0.05 (0.05)	0.59 (7.88)	90.8
9	-1.64 (-1.33)	-1.37 (-2.61)	-2.30 (-1.77)	-2.23 (-3.36)	1.34 (12.29)	0.91 (5.72)	0.17 (1.06)	0.24 (2.59)	-0.19 (-0.25)	0.19 (0.13)	0.64 (5.62)	79.4
10	-3.59 (-2.82)	-3.47 (-5.13)	-4.91 (-3.98)	-4.82 (-7.82)	1.01 (5.10)	0.93 (3.15)	0.05 (0.21)	0.14 (1.14)	0.48 (0.35)	-3.19 (-1.45)	0.72 (2.95)	52.7
1-10	3.66 (4.43)	3.69 (4.48)	4.86 (5.79)	4.85 (6.30)	0.02 (0.06)	-0.77 (-1.98)	0.08 (0.27)	-0.13 (-0.82)	-0.59 (-0.31)	1.34 (0.51)	-0.78 (-2.73)	11.3

Table 5.10: **Characteristics of (Un)Successful Option Traders**

This table reports time series averages of monthly cross-sectional averages of investor characteristics for decile portfolios of option investors formed on the basis of past one-year net return. As an exception, the numbers we report for account value and income are time series averages of monthly cross-sectional median values. Decile 1 contains the 10% of investors with the highest return during the ranking period and decile 10 includes the worst 10% performers in the ranking period. *Value* is the market value of all assets in the investor's account. *Turnover* is the average value of purchases and sales in a given month divided by beginning-of-the-month account value. *Men* is the percentage of accounts in a given decile portfolio held by men and *Age* is the age of the primary accountholder. *Income* is the monthly gross income assigned to investors based on their zip-code. *Experience* is the number of months the investor has been trading.

Decile	Value (€)	Turnover (%)	Men (%)	Age (years)	Income (€)	Experience (months)
1	26,341	15.9	70.6	44.1	2,706	41.2
2	25,604	5.9	64.5	46.9	2,611	41.8
3	27,023	3.7	71.2	47.0	2,714	42.0
4	18,853	5.3	67.1	49.5	2,473	41.8
5	12,401	8.1	73.7	44.9	2,552	41.7
6	14,001	8.7	76.3	54.1	2,498	41.9
7	4,261	9.6	84.1	42.1	2,316	41.0
8	5,382	8.4	83.0	46.8	2,323	40.7
9	3,235	16.0	84.6	45.3	2,404	39.5
10	1,813	76.7	85.1	48.1	2,279	34.8
1-10	23,106	-60.8	-14.5	4.0	427	6.4

Despite the poor performance of the average option trader, we do identify a small group of investors who consistently manage to outperform the others. Option traders who are in the top decile portfolio based on past one-year performance continue to outperform investors in the bottom decile over the next year. We further show that persistence in trading costs explains only part of total performance persistence. The bottom deciles tend to consist of male investors with little experience and low income who hold small accounts with high turnover. These results suggest that most option traders lose money due to excessive trading and a lack of knowledge. We conclude that trading hurts investor performance and that trading options hurts most.

5.A Appendix: Cross-Sectional Risk and Style Adjustment

In this appendix we show how to obtain an unbiased estimate of the relation between alphas and characteristics without estimating a first-stage time series regression for every investor. Each month, we estimate a cross-sectional regression of returns on investor characteristics. We then regress the vector of monthly cross-sectional coefficients γ_{lt} for each of the l investor characteristics on a constant and the time series of risk and style factor realizations F_{kt} ,

$$\gamma_{lt} = \delta_{l0} + \sum_{k=1}^K \delta_{lk} F_{kt} + \omega_{lt}. \quad (5.A.1)$$

To see that the intercept δ_{l0} in this time series regression is an unbiased estimate of the cross-sectional relation between alpha and investor characteristic l , consider the following. Suppose that portfolio returns are generated by the factor model

$$R_{it} = \alpha_i + \sum_{k=1}^K \beta_{ik} F_{kt} + \epsilon_{it}. \quad (5.A.2)$$

We are interested in the cross-sectional relation between α_i and Z_{ilt} , the value of characteristic l for investor i at time t . Denote the true coefficient vector of this relation between alphas and characteristics by δ . The cross-sectional regression of portfolio returns in month t on investor characteristics produces the following coefficient vector,

$$\gamma_t = (\mathbf{Z}'_t \mathbf{Z}_t)^{-1} \mathbf{Z}'_t \mathbf{R}_t = \delta + \mathbf{F}_t [(\mathbf{Z}'_t \mathbf{Z}_t)^{-1} \mathbf{Z}'_t \beta]. \quad (5.A.3)$$

Therefore, as noted by Brennan, Chordia, and Subrahmanyam (1998), the intercept δ_{l0} in the time series regression of the monthly cross-sectional parameter estimates γ_{lt} on the vector of factor realizations $\mathbf{F}_t$ is an unbiased estimate of δ_l if the factor premia are serially uncorrelated. Intuitively, the time series regression (5.A.1) purges the cross-sectional coefficients of their factor dependent component. The standard error of δ_{l0} is the standard error of the purged estimator. The coefficients on the factor premia $\mathbf{F}_t$ are unbiased estimates of $(\mathbf{Z}'_t \mathbf{Z}_t)^{-1} \mathbf{Z}'_t \beta$ and measure the relation between factor loadings and investor characteristics.

Chapter 6

Conclusion

This thesis sheds new light on the dynamic behavior of risk in financial markets and the behavior of individual investors who trade in these markets. A better understanding of volatility dynamics is important for both academics and practitioners because most financial decisions are based upon the tradeoff between risk and return. The empirical evidence that risk varies through time is overwhelming. A quick look at a stock chart immediately reveals distinct periods of high and low volatility. However, some basic questions are still left unanswered. First, it is not completely clear why the volatility of markets and individual stocks changes over time. Second, because the exact sources of time variation in risk are unknown, another unresolved issue is how to properly model these risk dynamics. Third, the ultimate open question is how fluctuations in risk affect the pricing of securities. In addition, because risk plays such an important role in investment decisions, the question arises how individual investors, who are generally considered to be less sophisticated than professional investors, respond to movements in stock markets. This chapter summarizes the answers to these fundamental questions based on the empirical research presented in this dissertation. Furthermore, it discusses the broader implications of these findings and provides avenues for future research.

6.1 Summary of Main Findings

Chapter 2 shows that variables that predict time variation in expected returns also pick up pronounced fluctuations in the exposures of portfolios sorted on firm size and book-to-market equity to the three risk factors in the Fama and French (1993) model. Cross-sectional variation in these betas is driven by the two firm characteristics size and book-to-market whereas time variation in the betas depends on the default spread, which proxies for macroeconomic conditions. Interaction terms between the firm characteristics and the default spread allow the relation between these characteristics and betas to vary with the business cycle. This captures the effect that value firms become more risky relative to growth firms in bad economic times when the price of risk is high, because they are less flexible in adjusting to

worsening economic conditions. Although a parametric specification of the relation between beta and firm-specific and macroeconomic variables is appealing from a theoretical point of view, an important practical problem with this approach is that the investor's set of conditioning information is unknown. Including too many variables makes it difficult to estimate parameters with any degree of precision but misspecification of the set of conditioning variables can lead to omitted-variable bias in the estimates of conditional asset pricing models. Hence, the critical assumption is that the relevant conditioning information is well summarized by a few variables.

Chapter 3 addresses this problem by combining a parametric specification of conditional betas with a nonparametric approach based on data-driven filters. Since the nonparametric approach does not explicitly link variation in risk to economic state variables, it is more robust to misspecification of beta dynamics. In addition, because it uses high-frequency data over a pre-specified time window with a set of optimally chosen weights, the precision and timeliness of the realized beta estimates is increased. The empirical analysis in chapter 3 further shows that the preferred combination of fundamental and realized betas varies across stocks and over time. For the average firm, fundamental and realized betas are equally important but for some stocks the set of conditioning variables underlying the fundamental beta better describes the evolution of beta. For other firms, however, movements in beta are not adequately captured by this set of firm-level and macroeconomic variables and more weight is given to the realized beta estimate. The optimal mix also varies over time because the fundamental beta picks up long-term changes in beta driven by structural changes in the economic environment and in firm-specific conditions. In contrast, the realized beta can capture short-term fluctuations in beta in periods of high market volatility as it is based on high-frequency data.

Chapter 4 studies long- and short-term risk dynamics at the aggregate level. Changes in market variance are broken down into changes in the variances of individual stocks and changes in the correlations between these stocks. Subsequently, market volatility, idiosyncratic volatilities, and correlations are decomposed into long and short run components. High-frequency return volatility is modeled as an asymmetric unit GARCH process that captures the stylized fact that negative stock returns have a larger impact on volatility than positive returns. The low-frequency component of volatility is modeled using an exponential spline with equally spaced knots. This approach to modeling long run volatility is completely data-driven and therefore less prone to misspecification. Chapter 4 also separates high- and low-frequency parts of conditional correlations by taking a factor model approach that combines the high- and low-frequency components of systematic and idiosyncratic risk with idiosyncratic correlations from a dynamic conditional correlation (DCC) model. This method also captures the influence of latent variables and time-varying betas on conditional correlations. The results show that long-term volatilities and correlations display strong counter-cyclical behavior, which highlights the importance of allowing unconditional volatilities and correlations to vary over time.

The risk dynamics documented in this thesis have important implications for asset pricing. Chapter 2 shows that allowing for time-varying risk loadings in the Fama and French (1993) three-factor model improves the performance of the model in explaining time variation in expected returns on European size-B/M sorted portfolios. However, for some portfolios pricing errors remain large and predictable. Furthermore, the conditional three-factor model does not capture the explanatory power of momentum variables for the cross-section of portfolio returns. In chapter 3 the focus is on individual stocks, whose betas exhibit stronger time variation than those of portfolios because most risk dynamics at the individual stock level cancel out at the portfolio level. The main reason for aggregating stocks into portfolios is to reduce measurement error in beta estimates. The crucial assumption of the portfolio approach is that all stocks that are grouped together in a given portfolio share the same risk characteristics. Consequently, if firms in a particular portfolio have different exposures to other determinants of risk than the characteristics they are sorted on, this method leads to errors when stocks are assigned the beta of the portfolio they belong to.

Chapter 3 provides strong evidence of cross-sectional variation in firm-specific betas within each of the 25 size-B/M portfolios, which implies that the homogeneity assumption underlying the portfolio approach is violated. This finding has important consequences for cross-sectional tests of asset pricing models. In particular, sorting stocks into portfolios leads to a loss of information contained in individual stocks because it reduces the cross-sectional dispersion in beta. As a result, using individual stocks as test assets instead of portfolios yields more precise estimates of factor risk premia. In chapter 3 a Bayesian panel data model is proposed that exploits the information in the cross-section of stocks to estimate the betas of individual stocks more precisely. Hierarchical prior distributions are specified that impose a common structure on the model parameters while still allowing for firm-level heterogeneity. Because only the parameters of these common distributions need to be estimated, the firm-specific beta estimates from the panel regression are more precise than those obtained from estimating a separate time series regression for every firm. The improved specification and more accurate estimation of time-varying betas lead to a sharp increase in the pricing ability of the conditional CAPM when individual stocks are used as test assets. The estimate of the market risk premium is positive and significant and is not affected by the inclusion of firm characteristics in the cross-sectional regressions. The time-varying betas from the panel model are also used to predict the covariance matrix of stock returns. Minimum variance portfolios based on this covariance matrix are superior in terms of out-of-sample standard deviation to portfolios constructed using competing approaches to forecasting covariances.

Chapter 4 discusses the pricing of shocks to long and short run market risk, individual risks, and market-wide correlations. It is shown that the model-free implied stock market variance computed from index option data systematically exceeds the realized market variance calculated from high-frequency returns. This signals the existence of a variance risk premium in index options, which means that the risk that market variance changes is priced. The decomposition of mar-

ket variance into individual variances and correlations shows that variance risk is priced in index options because correlation risk is priced. Variance and correlation risk premia increase sharply during periods of market turmoil, such as the LTCM and Asian crises. It turns out that the market variance risk premium is a strong predictor of monthly and quarterly stock market returns even after controlling for traditional return predictors. The empirical results further show that the predictive power of the market variance premium is completely driven by the correlation risk premium. The average individual variance risk premium is negative and does not predict returns. Chapter 4 also shows that the price of market volatility risk is significantly negative in the cross-section of stock returns, because investors prefer stocks that pay off after an unexpected increase in aggregate uncertainty. Innovations in long- and short-term market volatility both carry a significant price of risk, which indicates that volatility risk is priced at different horizons. Long-term market volatility risk is priced because shocks to average long-term idiosyncratic volatility and average long-term correlations are priced. In contrast, innovations in short run market volatility are only priced because short-term correlation risk is priced. The price of correlation risk is negative because investors care about shocks to market-wide correlations that reduce the benefits from diversification. Importantly, the variance and correlation factors have explanatory power for the cross-section of returns beyond the explanatory power of traditional risk factors, which suggests that they capture different sources of systematic risk.

Chapter 5 focuses on the behavior and performance of individual investors in option markets and examines whether these traders act rationally and understand the risk and return characteristics of these securities. Because the value of options directly depends on the volatility of the underlying asset, they are very suitable for transferring risks. Investors can use options to hedge the downward risk of their stock portfolio but can also employ option strategies to speculate on directional price movements or changes in volatility. Chapter 5 shows that individual investors who trade options significantly underperform those who only trade common stocks. The losses of option traders can be attributed to high transaction costs and to overreaction to past stock market movements. In particular, after the burst of the tech bubble, option investors took positions to speculate on a further market decline. As a result, they missed a substantial part of the market recovery and performed poorly compared to other investors. A detailed analysis of the trades and holdings of option investors reveals that the majority of investors trade out-of-the money options without holding the underlying stock. This suggests that most individual investors use options to increase rather than decrease their risks. Further evidence that investors trade options for gambling purposes is provided by linking option trading to a number of investor characteristics. It is found that single men with low income and little investment experience are more prone to trade options. It is known that this group of people also have a higher propensity to participate in lotteries. Although the average option investor performs poorly, chapter 5 identifies a small number of option traders who consistently outperform all other investors. Successful option investors trade less, hold larger accounts, are more experienced, and are more likely to be female than unsuccessful traders.

6.2 Implications of Empirical Results

The insights gained from the research presented in this dissertation are useful for academics, investors, managers, and policymakers. First, for risk managers in financial institutions the results are important because the risk exposure of their portfolios usually has to stay within predetermined limits. The evidence of time-varying risk loadings in chapters 2 and 3 implies that managers should constantly monitor the risks of their portfolios and adjust their holdings if necessary. In addition, the finding in chapter 4 that sudden changes in long and short run variances and correlations are priced risks in stock markets means that risk managers should also measure the exposure to these risks. For instance, Buraschi, Kosowski, and Trojani (2009) show that a hedge fund's ability to enter long-short positions can reduce its market beta but exposes the fund to unexpected changes in correlations. Accounting for long- and short-term fluctuations in variances and correlations is also essential for value-at-risk calculations.

Second, precise estimates of risk are important for performance evaluation. In particular, if the asset pricing model used to measure abnormal performance does not account for time-varying risk loadings or if it excludes important risk factors like the variance and correlation factors identified in this thesis, fund managers can generate positive abnormal returns without genuine investment skills.

Third, the implementation of portfolio theory requires estimates of expected returns and covariances, which are hard to measure when the investment universe is large. The Bayesian panel model proposed in chapter 3 can be used to increase the precision of beta forecasts of individual stocks and, consequently, the accuracy of forecasts of the covariance matrix of stock returns, which improves out-of-sample portfolio performance. Furthermore, an important consequence of the finding in chapter 4 that investors care about shocks to aggregate uncertainty and market-wide correlations, is that efficient portfolios should maximize expected returns given their return variances and the covariances of their returns with innovations in variances and correlations. Specifically, Buraschi, Porchia, and Trojani (2009) show that when the covariance matrix of returns is stochastic, the optimal portfolio for an investor with constant relative risk aversion preferences includes a significant correlation hedging component.

Fourth, an accurate estimate of a company's beta is an important prerequisite for successful capital budgeting decisions because the expected future cash flows of a project must be discounted at the firm's risk-adjusted rate of return. Imprecise estimates of risk lead to flawed estimates of the cost of capital and inaccurate valuations of projects. Because the Bayesian panel model delivers precise estimates of firm-level betas even when little return data is available, it can already be used by managers shortly after the initial public offering (IPO) of their firm.

Fifth, betas are often used to calculate abnormal returns in event studies, which measure the effect of mergers and acquisitions, earnings announcements, IPOs, and other corporate events on stock returns. Since these events are usually firm-specific, precise measures of individual stock betas are required, which can again be obtained from the Bayesian model in chapter 3. That chapter also stresses

the need to use individual stocks in cross-sectional tests of asset pricing models, because this leads to more precise estimates of risk premia and pricing errors, thereby increasing the power of the tests.

Finally, the finding in chapter 5 that clients of an online broker who trade options lose large amounts of money has implications for policymakers. This issue is especially important as the results indicate that investors with the lowest income trade most and also lose most. In addition, Anderson (2008) shows for a sample of Swedish investors that for those who trade most, the value of their online portfolio constitutes a large share of their total investments in risky assets. He concludes that “trading losses are mainly carried by those who can afford them the least.” The irrational behavior investors exhibit in trading options also casts doubts on their ability to make other financial decisions, like saving and investing for their retirement. Campbell (2006) shows that poorer and less educated households are more likely to make investment mistakes, including nonparticipation in risky asset markets and the failure to diversify their portfolios. To prevent people from making these mistakes, governments should seek to improve their financial literacy by setting up financial education programs. Thaler and Sunstein (2008) argue that a careful choice of default options, increased disclosure, and better product design can also help people in making sound financial decisions.

6.3 Suggestions for Future Research

This dissertation has addressed several important questions related to the measurement and pricing of time-varying risk in stock markets and the behavior and performance of individual investors in option markets. It also raises a number of new questions that should be addressed in future research.

First, more work should be done to identify the underlying economic sources of the long and short run variance and correlation risks discussed in chapter 4. Long- and short-term volatilities and correlations can be modeled using a non-parametric approach and linked to business cycle variables *ex post*, as in Rangel and Engle (2009). Another promising approach is to directly model the influence of macroeconomic factors on the covariance matrix of returns. Engle, Ghysels, and Sohn (2009) set up a component model in which variances are driven by economic sources. It would be interesting to extend this framework to correlations. To mitigate potential biases from misspecifying these risk dynamics, a combination of parametric and nonparametric approaches can be used, similar to the mixed beta specification in chapter 3 of this thesis. Another natural extension of the analyses in chapters 3 and 4 is to model long- and short-term movements in market betas and study the pricing of both beta components in the cross-section of returns.

Second, structural models should be developed to explain the wedge between implied and realized variances and correlations. Drechsler and Yaron (2008) argue that changes in the market variance risk premium reflect changes in agents’ perceptions of the risk of jumps in long run consumption growth and in the volatility of consumption. Related to this is a recent strand of literature that tries to ex-

plain the high equity premium and high stock market volatility using time-varying risk of rare disasters (Gabaix (2009) and Wachter (2009)). Time-varying severity of rare disasters increases stock market volatility because it induces changes in discount rates and, consequently, changes in stock prices. An alternative source of variation in variances and variance risk premia is time-varying risk aversion. A first step in this direction is taken by Bekaert and Engstrom (2009), who set up a structural model in which the consumption growth process is hit by good and bad shocks and where agents have preferences as in the habit formation model of Campbell and Cochrane (1999). In this framework, the variance risk premium increases with risk aversion and with shocks that generate negative skewness of consumption growth. Buraschi, Trojani, and Vedolin (2009) allow for heterogeneity in beliefs and argue that a common earnings disagreement component explains the dynamics of variance and correlation risk premia. Exploring the cross-sectional implications of these models can give further insights in the role of variance and correlation risk in stock markets and explain why firms differ in their exposure to long- and short-term variance and correlation factors.

Third, it is not clear whether the irrational behavior of individual investors documented in this thesis has a distorting impact on option prices. The effect of these traders on prices mainly depends on their market impact and on the behavior of other investors. Han (2008) argues that the sentiment of institutional investors explains time variation in the slope of the volatility smile in index options. However, Lakonishok, Lee, Pearson, and Poteshman (2007) show that sophisticated option investors did not exhibit irrational behavior during the stock market bubble at the end of the 1990s. Goyal and Saretto (2009) argue that the variance risk premium reflects option mispricing, driven by overreaction to current stock returns that leads to misestimation of future volatility. It would be interesting to directly test this hypothesis using data on the option trading activity of both individual and institutional investors.

Fourth, the finding in chapter 5 that most individual investors do not sufficiently understand the risk and return characteristics of options, raises the question whether investors learn from their mistakes and whether education can improve their performance. Initial evidence of Seru, Shumway, and Stoffman (2009) suggests that inexperienced investors do not learn fast enough to prevent behavioral biases from affecting asset prices. Moreover, an important part of their actual learning occurs when they cease trading after realizing that their inherent ability is poor rather than continuing to trade and improving their skills over time. A useful extension of this research is to compare the effect of trading experience on investor performance to the influence of financial education.

I close this chapter with a quote from Fischer Black. Although the option pricing model he developed with Scholes assumes that volatility is constant, Black (1976) notes that in reality “nothing is really constant. Volatilities themselves are not constant, and we cannot write down the process by which the volatilities change with any assurance that the process itself will stay fixed. We will have to keep updating our description of the process.” I hope that the research presented in this dissertation has updated and extended our knowledge of these risk dynamics.

Bibliography

- Adrian, T., and F. Franzoni, 2009, Learning about beta: Time-varying factor loadings, expected returns, and the conditional CAPM, *Journal of Empirical Finance* 16, 537–556.
- Adrian, T., and J. Rosenberg, 2008, Stock returns and volatility: Pricing the long-run and short-run components of market risk, *Journal of Finance* 63, 2997–3030.
- Agarwal, V., and N. Naik, 2004, Risks and portfolio decisions involving hedge funds, *Review of Financial Studies* 17, 63–98.
- Amihud, Y., 2002, Illiquidity and stock returns: Cross-section and time-series effects, *Journal of Financial Markets* 5, 31–56.
- Ammann, M., and M. Verhofen, 2008, Testing conditional asset pricing models using a markov chain Monte Carlo approach, *European Financial Management* 14, 391–418.
- Andersen, T., T. Bollerslev, F. Diebold, and H. Ebens, 2001, The distribution of realized stock return volatility, *Journal of Financial Economics* 61, 43–76.
- Andersen, T., T. Bollerslev, F. Diebold, and J. Wu, 2005, A framework for exploring the macroeconomic determinants of systematic risk, *American Economic Review* 95, 398–404.
- Anderson, A., 2008, Is online trading gambling with peanuts?, Working Paper, Stockholm University.
- Ang, A., and J. Chen, 2007, CAPM over the long run: 1926–2001, *Journal of Empirical Finance* 14, 1–40.
- Ang, A., R. Hodrick, Y. Xing, and X. Zhang, 2006, The cross-section of volatility and expected returns, *Journal of Finance* 61, 259–299.
- Ang, A., and D. Kristensen, 2009, Testing conditional factor models, Working Paper, Columbia University.

BIBLIOGRAPHY

- Ang, A., J. Liu, and K. Schwarz, 2008, Using stocks or portfolios in tests of factor models, Working Paper, Columbia University.
- Avramov, D., and T. Chordia, 2006a, Asset pricing models and financial market anomalies, *Review of Financial Studies* 19, 1001–1040.
- Avramov, D., and T. Chordia, 2006b, Predicting stock returns, *Journal of Financial Economics* 82, 387–415.
- Bali, T., 2008, The intertemporal relation between expected return and risk, *Journal of Financial Economics* 87, 101–131.
- Bali, T., N. Cakici, X. Yan, and Z. Zhang, 2005, Does idiosyncratic risk really matter?, *Journal of Finance* 60, 905–929.
- Bali, T., and R. Engle, 2008, Investigating ICAPM with dynamic conditional correlations, Working Paper, New York University.
- Bansal, R., and A. Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.
- Banz, R., 1981, The relation between return and market value of common stocks, *Journal of Financial Economics* 9, 3–18.
- Barber, B., and T. Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *Journal of Finance* 55, 773–806.
- Barber, B., and T. Odean, 2001, Boys will be boys: Gender, overconfidence, and common stock investment, *Quarterly Journal of Economics* 116, 261–292.
- Barber, B., and T. Odean, 2002, Online investors: Do the slow die first?, *Review of Financial Studies* 15, 455–487.
- Barberis, N., M. Huang, and T. Santos, 2001, Prospect theory and asset prices, *Quarterly Journal of Economics* 116, 1–53.
- Basu, S., 1977, Investment performance of common stocks in relation to their price-earnings ratios: A test of the efficient market hypothesis, *Journal of Finance* 12, 129–156.
- Bauer, R., M. Cosemans, and P. Eichholtz, 2009, Option trading and individual investor performance, *Journal of Banking and Finance* 33, 731–746.
- Bauer, R., M. Cosemans, and P. Schotman, 2009, Conditional asset pricing and stock market anomalies in Europe, *European Financial Management*, Forthcoming.
- Bauwens, L., M. Lubrano, and J. Richard, 1999, *Bayesian Inference in Dynamic Econometric Models* (Oxford University Press: Oxford).

- Bekaert, G., and E. Engstrom, 2009, Asset return dynamics under bad environment-good environment fundamentals, Working Paper, Columbia University.
- Bekaert, G., E. Engstrom, and Y. Xing, 2009, Risk, uncertainty and asset prices, *Journal of Financial Economics* 91, 59–82.
- Bekaert, G., R. J. Hodrick, and X. Zhang, 2009a, International stock return comovements, *Journal of Finance*, Forthcoming.
- Bekaert, G., R. J. Hodrick, and X. Zhang, 2009b, Is there a trend in idiosyncratic volatility?, Working Paper, Columbia University.
- Berglund, T., and J. Knif, 1999, Accounting for the accuracy of beta estimates in CAPM tests on assets with time-varying risks, *European Financial Management* 5, 29–42.
- Berk, J., R. Green, and V. Naik, 1999, Optimal investment, growth options, and security returns, *Journal of Finance* 54, 1553–1607.
- Bernstein, P., 1996, *Against the Gods: The Remarkable Story of Risk* (Wiley: New York).
- Black, F., 1976, Studies of stock price volatility changes, in *Proceedings of the 1976 Meetings of the Business and Economic Statistics Section, American Statistical Association* pp. 177–181.
- Black, F., M. Jensen, and M. Scholes, 1972, The capital asset pricing model: Some empirical tests, in M. Jensen, ed.: *Studies in the Theory of Capital Markets* (Praeger: New York).
- Black, F., and M. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.
- Bollerslev, T., R. Engle, and D. Nelson, 1994, ARCH models, in R. Engle, and D. McFadden, ed.: *Handbook of Econometrics* . pp. 2959–3038 (North-Holland: Amsterdam).
- Bollerslev, T., R. Engle, and J. Wooldridge, 1988, A capital asset pricing model with time-varying covariances, *Journal of Political Economy* 96, 116–131.
- Bollerslev, T., M. Gibson, and H. Zhou, 2009, Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities, *Journal of Econometrics*, Forthcoming.
- Bollerslev, T., G. Tauchen, and H. Zhou, 2009, Expected stock returns and variance risk premia, *Review of Financial Studies* 22, 4463–4492.
- Bollerslev, T., and B. Zhang, 2003, Measuring and modeling systematic risk in factor pricing models using high-frequency data, *Journal of Empirical Finance* 10, 533–558.

BIBLIOGRAPHY

- Brennan, M., T. Chordia, and A. Subrahmanyam, 1998, Alternative factor specifications, security characteristics, and the cross-section of expected stock returns, *Journal of Financial Economics* 49, 345–373.
- Britten-Jones, M., and A. Neuberger, 2000, Option prices, implied price processes, and stochastic volatility, *Journal of Finance* 55, 839–866.
- Brown, G., and M. Cliff, 2004, Investor sentiment and the near-term stock market, *Journal of Empirical Finance* 11, 1–27.
- Brunnermeier, M., and S. Nagel, 2004, Hedge funds and the technology bubble, *Journal of Finance* 59, 2013–2038.
- Buraschi, A., R. Kosowski, and F. Trojani, 2009, When there is no place to hide: Correlation risk and the cross-section of hedge fund returns, Working Paper, Imperial College London.
- Buraschi, A., P. Porchia, and F. Trojani, 2009, Correlation risk and optimal portfolio choice, *Journal of Finance*, Forthcoming.
- Buraschi, A., F. Trojani, and A. Vedolin, 2009, When uncertainty blows in the orchard: Comovement and equilibrium volatility risk premia, Working Paper, Imperial College London.
- Busse, J., and P. Irvine, 2006, Bayesian alphas and mutual fund persistence, *Journal of Finance* 61, 2251–2288.
- Campbell, J., 1991, A variance decomposition for stock returns, *Economic Journal* 101, 157–179.
- Campbell, J., 2006, Household finance, *Journal of Finance* 61, 1553–1604.
- Campbell, J., and J. Cochrane, 1999, By force of habit: A consumption based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Campbell, J., M. Lettau, B. Malkiel, and Y. Xu, 2001, Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk, *Journal of Finance* 56, 1–43.
- Campbell, J., A. Lo, and A. MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press: Princeton).
- Campbell, J., and R. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.
- Carhart, M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.
- Carr, P., and L. Wu, 2009, Variance risk premiums, *Review of Financial Studies* 22, 1311–1341.

- Chan, L., J. Karceski, and J. Lakonishok, 1999, On portfolio optimization: Forecasting covariances and choosing the risk model, *Review of Financial Studies* 12, 937–974.
- Chang, B., P. Christoffersen, K. Jacobs, and G. Vainberg, 2009, Option-implied measures of equity risk, Working Paper, McGill University.
- Chernov, M., R. Gallant, E. Ghysels, and G. Tauchen, 2003, Alternative models for stock price dynamics, *Journal of Econometrics* 116, 225–257.
- Chib, S., and E. Greenberg, 1995, Understanding the Metropolis-Hastings algorithm, *The American Statistician* 49, 327–335.
- Chordia, T., and L. Shivakumar, 2002, Momentum, business cycle, and time-varying expected returns, *Journal of Finance* 57, 985–1019.
- Cochrane, J., 2005, *Asset Pricing* (Princeton University Press: Princeton) 2nd edn.
- Colacito, R., R. Engle, and E. Ghysels, 2009, A component model for dynamic correlations, Working Paper, University of North Carolina at Chapel Hill.
- Cosemans, M., R. Frehen, P. Schotman, and R. Bauer, 2009, Efficient estimation of firm-specific betas and its benefits for asset pricing tests and portfolio choice, Working Paper, Maastricht University.
- Coval, J., D. Hirshleifer, and T. Shumway, 2005, Can individual investors beat the market?, Working Paper, Harvard University.
- Coval, J., and T. Shumway, 2005, Do behavioral biases affect asset prices?, *Journal of Finance* 60, 1–34.
- Cremers, M., 2006, Multifactor efficiency and Bayesian inference, *Journal of Business* 79, 2951–2998.
- Daniel, K., and S. Titman, 1997, Evidence on the characteristics of cross sectional variation in stock returns', *Journal of Finance* 52, 1–33.
- DeBondt, W., and R. Thaler, 1985, Does the stock market overreact?, *Journal of Finance* 40, 793–805.
- DeMiguel, V., L. Garlappi, and R. Uppal, 2009, Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?, *Review of Financial Studies* 22, 1915–1953.
- Dorn, D., and P. Sengmueller, 2009, Trading as entertainment, *Management Science* 55, 591–603.
- Drechsler, I., and A. Yaron, 2008, What's vol got to do with it, Working Paper, University of Pennsylvania.

BIBLIOGRAPHY

- Driessen, J., P. Maenhout, and G. Vilkov, 2009, The price of correlation risk: Evidence from equity options, *Journal of Finance* 64, 1377–1406.
- Eiling, E., and B. Gerard, 2007, Dispersion, equity returns correlations and market integration, Working Paper, University of Toronto.
- Engle, R., 2002, Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models, *Journal of Business and Economic Statistics* 20, 339–350.
- Engle, R., 2009a, *Anticipating Correlations: A New Paradigm for Risk Management* (Princeton University Press: Princeton).
- Engle, R., 2009b, High dimension dynamic correlations, in J. Castle, and N. Shephard, ed.: *The Methodology and Practice of Econometrics: A Festschrift in Honour of David F. Hendry* . pp. 122–148 (Oxford University Press: Oxford).
- Engle, R., and R. Colacito, 2006, Testing and valuing dynamic correlations for asset allocation, *Journal of Business and Economic Statistics* 24, 238–253.
- Engle, R., E. Ghysels, and B. Sohn, 2009, Stock market volatility and macroeconomic fundamentals, Working Paper, New York University.
- Engle, R., and G. Lee, 1999, A long run and short run component model of stock return volatility, in R. Engle, and H. White, ed.: *Cointegration, Causality, and Forecasting: A Festschrift in Honour of Clive W. J. Granger* . pp. 475–497 (Oxford University Press: Oxford).
- Engle, R., and J. Rangel, 2008, The spline-GARCH model for low-frequency volatility and its global macroeconomic causes, *Review of Financial Studies* 21, 1187–1222.
- Fama, E., 1965, The behavior of stock-market prices, *Journal of Business* 38, 34–105.
- Fama, E., and K. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.
- Fama, E., and K. French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427–465.
- Fama, E., and K. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Fama, E., and K. French, 1996, Multifactor explanations of asset pricing anomalies, *Journal of Finance* 51, 55–84.
- Fama, E., and K. French, 1997, Industry costs of equity, *Journal of Financial Economics* 43, 153–193.

- Fama, E., and K. French, 1998, Value versus growth: The international evidence, *Journal of Finance* 53, 1975–1999.
- Fama, E., and K. French, 2006, The value premium and the CAPM, *Journal of Finance* 61, 2163–2186.
- Fama, E., and K. French, 2008, Dissecting anomalies, *Journal of Finance* 63, 1653–1678.
- Fama, E., and J. MacBeth, 1973, Risk, return and equilibrium: Empirical tests, *Journal of Political Economy* 71, 607–636.
- Fama, E., and G. Schwert, 1977, Asset returns and inflation, *Journal of Financial Economics* 5, 115–146.
- Ferson, W., and C. Harvey, 1998, Fundamental determinants of national equity market returns: A perspective on conditional asset pricing, *Journal of Banking and Finance* 21, 1625–1665.
- Ferson, W., and C. Harvey, 1999, Conditioning variables and the cross section of stock returns, *Journal of Finance* 54, 1325–1360.
- Ferson, W., and R. Schadt, 1996, Measuring fund strategy and performance in changing economic conditions, *Journal of Finance* 51, 425–462.
- Franzoni, F., 2008, The changing nature of market risk, Working Paper, University of Lugano.
- French, K., and J. Poterba, 1991, Investor diversification and international equity markets, *American Economic Review* 81, 222–226.
- Fu, F., 2009, Idiosyncratic risk and the cross-section of expected stock returns, *Journal of Financial Economics* 91, 24–37.
- Gabaix, X., 2009, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, Working Paper, New York University.
- Gelman, A., J. Carlin, H. Stern, and D. Rubin, 2004, *Bayesian Data Analysis* (Chapman and Hall/CRC: Boca Raton, Florida) 2nd edn.
- Geweke, J., 2005, *Contemporary Bayesian Econometrics and Statistics* (Wiley: New York).
- Ghysels, E., 1998, On stable factor structures in the pricing of risk: Do time-varying betas help or hurt?, *Journal of Finance* 53, 549–573.
- Ghysels, E., P. Santa-Clara, and R. Valkanov, 2005, There is a risk-return trade-off after all, *Journal of Financial Economics* 76, 509–548.
- Ghysels, T., A. Harvey, and E. Renault, 1996, Stochastic volatility, in C. Rao, and G. Maddala, ed.: *Handbook of Statistics* . pp. 119–191 (North-Holland: Amsterdam).

BIBLIOGRAPHY

- Gibbons, M., S. Ross, and J. Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 57, 1121–1152.
- Glosten, L., and R. Jagannathan, 1994, A contingent claim approach to performance evaluation, *Journal of Empirical Finance* 1, 133–160.
- Goetzmann, W., A. Watanabe, and M. Watanabe, 2009, Investor expectations, business conditions, and the pricing of beta-instability risk, Working Paper, Yale School of Management.
- Gomes, J., L. Kogan, and L. Zhang, 2003, Equilibrium cross-section of returns, *Journal of Political Economy* 111, 693–732.
- Goyal, A., and P. Santa-Clara, 2003, Idiosyncratic risk matters!, *Journal of Finance* 58, 975–1007.
- Goyal, A., and A. Saretto, 2009, Cross-section of option returns and volatility, *Journal of Financial Economics* 94, 310–326.
- Griffin, J., X. Ji, and S. Martin, 2003, Momentum investing and business cycle risk: Evidence from pole to pole, *Journal of Finance* 58, 2515–2547.
- Grinblatt, M., and M. Keloharju, 2009, Sensation seeking, overconfidence, and trading activity, *Journal of Finance* 64, 549–578.
- Hamilton, J., 1994, *Time Series Analysis* (Princeton University Press: Princeton).
- Han, B., 2008, Investor sentiment and option prices, *Review of Financial Studies* 21, 387–414.
- Hansen, L., and S. Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, *Econometrica* 55, 587–613.
- Harvey, C., and G. Zhou, 1990, Bayesian inference in asset pricing tests, *Journal of Financial Economics* 26, 221–254.
- Heckman, J., 1976, The common structure of statistical models of truncation, sample selection and limited dependent variables and a simple estimator for such models, *Annals of Economic and Social Measurement* 5, 475–492.
- Heston, S., K. G. Rouwenhorst, and R. Wessels, 1999, The role of beta and size in the cross-section of european stock returns, *European Financial Management* 5, 9–27.
- Hodrick, R., 1992, Dividend yields and expected stock returns: Application to financial markets, *Review of Financial Studies* 5, 357–386.
- Jagannathan, R., and T. Ma, 2003, Risk reduction in large portfolios: Why imposing the wrong constraints helps, *Journal of Finance* 58, 1651–1683.

- Jagannathan, R., and Z. Wang, 1996, The conditional CAPM and cross-section of expected returns, *Journal of Finance* 51, 3–53.
- Jegadeesh, N., and S. Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.
- Jiang, G., and Y. Tian, 2005, The model-free implied volatility and its information content, *Review of Financial Studies* 18, 1305–1342.
- Jostova, G., and A. Philipov, 2005, Bayesian analysis of stochastic betas, *Journal of Financial and Quantitative Analysis* 40, 747–778.
- Kahneman, D., and A. Tversky, 1974, Judgment under uncertainty: Heuristics and biases, *Science* 185, 1124–1131.
- Kahneman, D., and A. Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263–290.
- Kandel, S., R. McCulloch, and R. Stambaugh, 1995, Bayesian inference and portfolio efficiency, *Review of Financial Studies* 8, 1–53.
- Karlsson, A., and L. Norden, 2007, Home sweet home: Home bias and international diversification among individual investors, *Journal of Banking and Finance* 31, 317–333.
- Keim, D., and A. Madhavan, 1998, The cost of institutional equity trades, *Financial Analysts Journal* 54, 50–69.
- Keim, D., and R. Stambaugh, 1986, Predicting returns in the bond and stock markets, *Journal of Financial Economics* 17, 357–390.
- Keynes, J., 1937, The general theory of employment, *Quarterly Journal of Economics* 51, 209–223.
- Knight, F., 1921, *Risk, Uncertainty, and Profit* (Houghton Mifflin: Boston).
- Krishnan, C., R. Petkova, and P. Ritchken, 2009, Correlation risk, *Journal of Empirical Finance* 16, 353–367.
- Kumar, A., 2009a, Dynamic style preferences of individual investors and stock returns, *Journal of Financial and Quantitative Analysis* 44, 607–640.
- Kumar, A., 2009b, Who gambles in the stock market?, *Journal of Finance* 64, 1889–1933.
- Lakonishok, J., I. Lee, N. Pearson, and A. Poteshman, 2007, Option market activity, *Review of Financial Studies* 20, 813–857.
- Lakonishok, J., A. Shleifer, and R. Vishny, 1994, Contrarian investment, extrapolation, and risk, *Journal of Finance* pp. 1541–1578.

BIBLIOGRAPHY

- Ledoit, O., and M. Wolf, 2003, Improved estimation of the covariance matrix of stock returns with an application to portfolios selection, *Journal of Empirical Finance* 10, 603–621.
- Lettau, M., and S. Ludvigson, 2001a, Consumption, aggregate wealth, and expected stock returns, *Journal of Finance* 56, 815–849.
- Lettau, M., and S. Ludvigson, 2001b, Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying, *Journal of Political Economy* 109, 1238–1287.
- Lewellen, J., 1999, The time-series relations among expected return, risk, and book-to-market, *Journal of Financial Economics* 54, 5–43.
- Lewellen, J., 2002, Momentum and autocorrelation in stock returns, *Review of Financial Studies* 15, 533–563.
- Lewellen, J., and S. Nagel, 2006, The conditional CAPM does not explain asset-pricing anomalies, *Journal of Financial Economics* 82, 289–314.
- Lewellen, J., S. Nagel, and J. Shanken, 2009, A skeptical appraisal of asset-pricing tests, *Journal of Financial Economics*, Forthcoming.
- Lintner, J., 1965, The valuation of risky assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics* 47, 13–37.
- Lo, A., and A. MacKinlay, 1990, Data-snooping biases in tests of financial asset pricing models, *Review of Financial Studies* 3, 431–467.
- Lucas, R., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1419–1446.
- Mahani, R., and A. Poteszman, 2008, Overreaction to stock market news and misvaluation of stock prices by unsophisticated investors: Evidence from the option market, *Journal of Empirical Finance* 15, 635–655.
- Markowitz, H., 1952, Portfolio selection, *Journal of Finance* 7, 77–91.
- Merton, R., 1973, An intertemporal asset pricing model, *Econometrica* 41, 867–887.
- Merton, R., 1987, A simple model of capital market equilibrium with incomplete information, *Journal of Finance* 42, 483–510.
- Millward-Brown, 2006, The retail investor 2006: Developments in the market of retail investors in the Netherlands, Amsterdam.
- Moskowitz, T., 2003, An analysis of covariance risk and pricing anomalies, *Review of Financial Studies* 16, 417–457.

- Ng, L., and F. Wu, 2007, The trading behavior of institutions and individuals in Chinese equity markets, *Journal of Banking and Finance* 31, 2695–2710.
- Pastor, L., and P. Veronesi, 2003, Stock valuation and learning about profitability, *Journal of Finance* 58, 1749–1789.
- Petersen, M., 2009, Estimating standard errors in finance panel data sets: Comparing approaches, *Review of Financial Studies* 22, 435–480.
- Petkova, R., and L. Zhang, 2005, Is value riskier than growth?, *Journal of Financial Economics* 78, 187–202.
- Poteshman, A., 2001, Underreaction, overreaction, and increasing misreaction to information in the options market, *Journal of Finance* 56, 851–876.
- Poteshman, A., and V. Serbin, 2003, Clearly irrational financial market behavior: Evidence from the early exercise of exchange traded stock options, *Journal of Finance* 58, 37–70.
- Rangel, J., and R. Engle, 2009, The factor-spline-GARCH model for high and low frequency correlations, Working Paper, New York University.
- Rosenberg, J., and R. Engle, 2002, Empirical pricing kernels, *Journal of Financial Economics* 64, 341–372.
- Ross, S., 2005, *Neoclassical Finance* (Princeton University Press: Princeton).
- Rouwenhorst, K. G., 1998, International momentum strategies, *Journal of Finance* 53, 267–284.
- Samuelson, P., 1965, Proof that properly anticipated prices fluctuate randomly, *Industrial Management Review* 6, 41–49.
- Santos, T., and P. Veronesi, 2004, Conditional betas, Working Paper, University of Chicago.
- Schmitz, P., M. Glaser, and M. Weber, 2009, Individual investor sentiment and stock returns - what do we learn from warrant traders?, Working Paper, University of Mannheim.
- Scholes, M., and J. Williams, 1977, Estimating betas from nonsynchronous data, *Journal of Financial Economics* 5, 309–327.
- Seru, A., T. Shumway, and N. Stoffman, 2009, Learning by trading, *Review of Financial Studies*, Forthcoming.
- Shanken, J., 1987, A Bayesian approach to testing portfolio efficiency, *Journal of Financial Economics* 19, 195–215.
- Shanken, J., 1990, Intertemporal asset pricing: An empirical investigation, *Journal of Econometrics* 45, 99–120.

BIBLIOGRAPHY

- Shanken, J., 1992, On the estimation of beta-pricing models, *Review of Financial Studies* 5, 1–34.
- Shapira, Z., and I. Venezia, 2001, Patterns of behavior of professionally managed and independent investors, *Journal of Banking and Finance* 25, 1573–1587.
- Sharpe, W., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance* 19, 425–442.
- Shiller, R., 2005, *Irrational Exuberance* (Princeton University Press: Princeton).
- Shleifer, A., 2000, *Inefficient Markets: An Introduction to Behavioral Finance* (Oxford University Press: Oxford).
- Thaler, R., and C. Sunstein, 2008, *Nudge: Improving Decisions about Health, Wealth, and Happiness* (Yale University Press: New Haven).
- Todorov, V., 2009, Variance risk premium dynamics: The role of jumps, *Review of Financial Studies*, Forthcoming.
- Vasicek, O., 1973, A note on using cross-sectional information in Bayesian estimation of security betas, *Journal of Finance* 28, 1233–1239.
- Vuolteenaho, T., 2002, What drives firm-level stock returns?, *Journal of Finance* 57, 233–264.
- Wachter, J., 2009, Can time-varying risk of rare disasters explain aggregate stock market volatility?, Working Paper, University of Pennsylvania.
- Wu, X., 2002, A conditional multifactor analysis of return momentum, *Journal of Banking and Finance* 26, 1675–1696.
- Zhang, L., 2005, The value premium, *Journal of Finance* 60, 67–103.

Nederlandse Samenvatting

Het begrip risico speelt een belangrijke rol in ons leven en neemt verschillende vormen aan. Risico heeft invloed op beslissingen in tal van gebieden, zoals economie, techniek, geneeskunde en wetenschap. Hoewel de precieze betekenis van risico afhangt van de context, verwijst de term doorgaans naar de kans op een verlies of andere ongewenste uitkomst. Omdat risico een belangrijke invloed heeft op onze samenleving, is het nauwkeurig meten en beheersen van risico van cruciaal belang voor economische groei en technologische vooruitgang. Deze ingewikkelde taken worden nog gecompliceerder door het dynamische karakter van risico.

Dit proefschrift gaat over het meten en het prijzen van tijdsvariërend risico in aandelenmarkten. Ofschoon het belang van risico in kansspelen al eeuwen geleden werd onderkend, werd het concept niet toegepast op aandelenmarkten totdat Harry Markowitz zijn portefeuilletheorie ontwikkelde in de jaren '50 van de vorige eeuw, waarvoor hij de Nobelprijs voor de Economie ontving in 1990 (samen met William Sharpe en Merton Miller). Volgens Markowitz (1952) zullen beleggers alleen portefeuilles kiezen die een maximaal verwacht rendement opleveren voor een bepaalde mate van risico. Het belangrijkste element in deze portefeuilletheorie is het concept van diversificatie. Dit begrip houdt in dat het risico van een portefeuille verminderd kan worden door het spreiden van vermogen over verschillende aandelen, mits de prijzen van deze aandelen niet volledig met elkaar samenhangen. Het risico van een portefeuille, gemeten door de variantie van het rendement, hangt daarom niet alleen af van de varianties van de individuele aandelen in de portefeuille maar ook van de covarianties tussen deze aandelen.

De inzichten uit de portefeuilletheorie vormden de basis voor Sharpe (1964) en Lintner (1965) bij de ontwikkeling van een model om financiële activa te waarderen, het zogeheten *Capital Asset Pricing Model* (CAPM). Zij laten zien dat onder bepaalde voorwaarden alle beleggers dezelfde beleggingsmogelijkheden hebben en dezelfde portefeuille van risicovolle effecten zullen bezitten, die daarom de marktportefeuille moet zijn. Ze kunnen deze belegging in de marktportefeuille combineren met een risicoloze positie om hun gewenste mate van risico te bereiken. Als alle beleggers dezelfde risicovolle portefeuille bezitten, dan is het relevante risico van een individueel aandeel de bijdrage aan het risico van de marktportefeuille.

Deze bijdrage wordt de bèta van het aandeel genoemd en is gedefinieerd als de covariantie tussen het rendement op het aandeel en het rendement op de gehele aandelenmarkt, gedeeld door de variantie van de marktportefeuille. Bèta meet dus de gevoeligheid van een aandeel ten opzichte van variatie in het rendement op de markt. Volgens het CAPM worden beleggers alleen gecompenseerd voor het nemen van systematisch risico, gemeten door bèta, omdat het specifieke risico van een aandeel door diversificatie kan worden geëlimineerd.

De eerste empirische toetsen van het CAPM bevestigden de voorspelling van het model dat de cross-sectie van verwachte rendementen volledig verklaard wordt door bèta. Latere studies ontdekten echter patronen in rendementen die inconsistent zijn met de implicaties van het CAPM. Zo blijken de aandelen van kleine bedrijven een hoger rendement op te leveren dan die van grote ondernemingen. Dit verschil in rendement kan niet worden verklaard door het verschil in bèta tussen grote en kleine bedrijven. Diverse verklaringen zijn aangedragen voor het falen van het CAPM. Aanhangers van de rationele waarderingstheorie en de efficiënte markt hypothese beweren dat de afwijkende patronen in rendementen verklaard kunnen worden door risico. Anderen geloven echter dat afwijkingen in aandelenkoersen optreden omdat beleggers zich irrationeel gedragen bij het verwerken van nieuwe informatie. Volgens Cochrane (2005) draait deze discussie uiteindelijk om de vraag of een waarderingstheorie de wijze beschrijft waarop de wereld werkt of de wijze waarop de wereld zou moeten werken.

Volgens de rationele theorie is een mogelijke reden voor het onvermogen van het CAPM en andere waarderingsmodellen om de cross-sectie van rendementen te verklaren de belangrijke veronderstelling van deze modellen dat het risico van aandelen constant is. Deze aanname is niet erg plausibel gezien het dynamische karakter van onze economieën en financiële markten. Als bèta in werkelijkheid fluctueert, dan zijn statische modellen zoals het CAPM verkeerd gespecificeerd en zullen ze een onvolledige beschrijving van aandelenrendementen geven. Deze observatie heeft geleid tot de ontwikkeling van conditionele waarderingsmodellen waarin bèta's kunnen veranderen. Het empirisch bewijs voor deze modellen is echter nog verre van eenduidig. Een belangrijke reden voor de tegenstrijdige resultaten is dat het moeilijk is de bèta's van individuele aandelen nauwkeurig te schatten, vooral wanneer deze bèta's tijdsvariërend zijn en wanneer het aantal datapunten beperkt is. Een ander probleem is dat het niet duidelijk is hoe tijdsvariatie in bèta's gemodelleerd moet worden. Hoewel het mogelijk is om robuuste, niet-parametrische technieken te gebruiken om veranderingen in risico te modelleren, moet er bij deze methoden een afweging gemaakt worden tussen de actualiteit en de accuratesse van de bèta schattingen.

Conjunctuurschommelingen hebben niet alleen invloed op het risico van individuele aandelen, maar ook op het risico van de marktportefeuille zelf. Omdat een verandering in marktrisico de verhouding tussen risico en rendement beïnvloedt, kan het een geprijsde risicofactor zijn in het intertemporele CAPM (ICAPM) ontwikkeld door Merton (1973). Terwijl het CAPM veronderstelt dat beleggers alleen geïnteresseerd zijn in het vermogen dat hun portefeuille één periode later oplevert, houdt het ICAPM er rekening mee dat in werkelijkheid beleggers ook letten op de

mogelijkheden die ze hebben om dit vermogen te beleggen aan het einde van deze periode. In het ICAPM willen risico-averse beleggers zich daarom beschermen tegen veranderingen in variabelen die het toekomstige beleggingsklimaat negatief beïnvloeden. Ze zijn daarom bereid om lagere verwachte rendementen te accepteren op aandelen die het goed doen als het verwachte toekomstige marktrendement daalt of de verwachte toekomstige volatiliteit van de aandelenmarkt stijgt.

Omdat de portefeuilletheorie laat zien dat de variantie van de marktportefeuille afhangt van de varianties en paarsgewijze correlaties van alle aandelen in de markt, moeten veranderingen in marktrisico gedreven worden door veranderingen in deze individuele varianties en correlaties. Dit betekent dat als markt variantierisico geprijsd is in de cross-sectie van aandelenrendementen, individueel variantierisico, correlatierisico of beiden geprijsd moeten zijn. Omdat nieuws een ander effect kan hebben op aandelenkoersen op lange termijn dan op korte termijn, kunnen varianties en correlaties zowel korte als lange termijn fluctuaties vertonen. Lange en korte termijn variantie- en correlatiefactoren kunnen daarom aparte risico's vertegenwoordigen en verschillend geprijsd zijn.

De systematische afwijkingen van aandelenkoersen van de voorspellingen van de efficiënte markt hypothese hebben geleid tot het ontstaan van een alternatief perspectief op financiële markten in de jaren '80, bekend onder de naam gedragsfinanciering. Deze zienswijze schrijft afwijkende patronen in aandelenrendementen toe aan irrationaliteit van beleggers en is gebaseerd op twee bouwstenen, te weten beperkte arbitrage en gedragsafwijkingen. Door arbitragebeperkingen kan het risicovol en kostbaar zijn voor rationele arbitrageanten om strategieën te implementeren die de verstorende invloed van irrationele beleggers op prijzen ongedaan maken. Risico's ontstaan bijvoorbeeld als er geen perfecte substituten zijn voor de aandelen die verkeerd geprijsd zijn waardoor arbitrageanten geen risicoloze positie kunnen innemen. Andere factoren die arbitrage belemmeren en daardoor prijsafwijkingen laten bestaan zijn transactiekosten en restricties op short selling.

De tweede pijler van gedragsfinanciering, gedragsafwijkingen, verwijst naar afwijkingen in de wijze waarop mensen verwachtingen vormen en in de wijze waarop ze beslissingen nemen gebaseerd op deze verwachtingen. Omdat de wereld complex is en de cognitieve vaardigheden van mensen beperkt zijn, gebruiken ze vuistregels als leidraad bij het nemen van beslissingen. Kahneman en Tversky (1974) laten zien dat deze regels in sommige situaties heel nuttig kunnen zijn maar in andere gevallen kunnen leiden tot systematische inschattingsfouten. Zo kunnen kansberekeningen verstoord worden wanneer mensen verbanden leggen die in werkelijkheid niet bestaan omdat een gebeurtenis lijkt op hun beeld van een andere gebeurtenis. Afwijkingen in verwachtingen kunnen ook ontstaan als gevolg van overmoed, waardoor mensen de nauwkeurigheid van hun beoordelingen overschatten. Ofschoon deze psychologische afwijkingen alleen invloed hebben op de prijzen van aandelen en andere financiële activa als arbitrage beperkt is en als de strategieën van irrationele beleggers gecorrigeerd zijn, hebben ze een directe invloed op het handelsgedrag en de prestaties van beleggers. Barber en Odean (2000) beweren bijvoorbeeld dat beleggers door overmoed teveel handelen, wat leidt tot slechte prestaties als gevolg van hoge transactiekosten.

Dit proefschrift werpt nieuw licht op het dynamische gedrag van risico in financiële markten en het gedrag van de beleggers die in deze markten handelen. Een beter begrip van veranderingen in risico is van belang voor zowel de wetenschap als het bedrijfsleven omdat de meeste financiële beslissingen gebaseerd zijn op de afweging tussen risico en rendement. Een vluchtige blik op een grafiek met aandelenkoersen laat meteen zien dat er afzonderlijke periodes bestaan met hoge en lage volatiliteit. Een aantal kernvragen zijn echter nog steeds niet beantwoord. Allereerst is het niet helemaal duidelijk waardoor het risico van markten en individuele aandelen verandert. Omdat de precieze oorzaken van tijdsvariatie in risico onbekend zijn, is een tweede onbeantwoorde vraag hoe deze veranderingen in risico gemodelleerd moeten worden. Een derde onopgelost vraagstuk is hoe fluctuaties in risico de waardering van aandelen beïnvloeden. Omdat risico zo een belangrijke rol speelt bij beleggingsbeslissingen, rijst tot slot de vraag hoe particuliere beleggers, die doorgaans worden beschouwd als minder ontwikkeld dan professionele beleggers, reageren op koersschommelingen in aandelenmarkten.

Hoofdstuk 2 bestudeert de risicodynamiek van portefeuilles die bestaan uit aandelen uit 16 Europese landen. De resultaten laten zien dat variabelen die tijdsvariatie in verwachte rendementen voorspellen ook sterke fluctuaties oppikken in de gevoeligheid van aandelenportefeuilles ten opzichte van de drie risicofactoren in het waarderingsmodel van Fama en French (1993). Cross-sectionele variatie in deze bèta's wordt gedreven door de marktwaarde van bedrijven en door de verhouding van hun boekwaarde tot hun marktwaarde terwijl tijdsvariatie in de bèta's afhangt van de *default spread*. Hoewel deze parametrische specificatie van de relatie tussen bèta's en bedrijfsspecifieke en macro-economische variabelen aantrekkelijk is vanuit een theoretisch oogpunt, is een belangrijk praktisch probleem dat het onbekend is welke variabelen gebruikt moeten worden. Het toevoegen van teveel variabelen maakt het lastig om de parameters met enige precisie te schatten maar het uitsluiten van relevante variabelen kan leiden tot afwijkingen in de schattingen van conditionele waarderingsmodellen.

Hoofdstuk 3 pakt dit probleem aan door het combineren van een parametrische specificatie van tijdsvariërende bèta's met een niet-parametrische methode. Omdat de niet-parametrische benadering de variatie in risico niet expliciet relateert aan economische variabelen is deze aanpak robuuster voor misspecificatie. Daarnaast maakt deze methode gebruik van hoogfrequente data en wordt er een functie gespecificeerd die meer gewicht toekent aan de meest recente observaties. Hierdoor wordt de nauwkeurigheid en de relevantie van de bèta schattingen sterk vergroot. De empirische analyse in hoofdstuk 3 laat verder zien dat de optimale combinatie van fundamentele bèta's uit de parametrische specificatie en door data gedreven bèta's uit de niet-parametrische specificatie per aandeel verschilt en over de tijd verandert. De optimale mix varieert omdat de fundamentele bèta lange termijn veranderingen in bèta oppikt die worden gedreven door structurele veranderingen in het economische klimaat en in bedrijfsspecifieke omstandigheden. De volledig door data gedreven bèta daarentegen kan door het gebruik van hoogfrequente data korte termijn fluctuaties in bèta oppikken tijdens periodes van hoge marktvolatiliteit.

Hoofdstuk 4 bestudeert veranderingen in de korte en lange termijn risico's van de totale aandelenmarkt. Veranderingen in de variantie van de marktportefeuille worden onderverdeeld in veranderingen in de varianties van individuele aandelen en veranderingen in de correlaties tussen deze aandelen. Vervolgens worden marktvolatiliteit, individuele volatiliteiten en correlaties opgesplitst in lange en korte termijn componenten. Hoogfrequente volatiliteit volgt een asymmetrisch GARCH proces terwijl de laagfrequente component van volatiliteit wordt gemodelleerd met behulp van een exponentiële *spline* met gelijk verdeelde knooppunten. Deze methode om lange termijn volatiliteit te modelleren is volledig data gedreven en daardoor minder gevoelig voor misspecificatie. Hoofdstuk 4 onderscheidt ook hoog- en laagfrequente delen van conditionele correlaties door gebruik te maken van een factor model dat de hoog- en laagfrequente componenten van systematisch en idiosyncratisch risico combineert met idiosyncratische correlaties uit een dynamisch conditioneel correlatie (DCC) model. De resultaten laten zien dat varianties en correlaties sterk anticyclisch gedrag vertonen.

De risicoschommelingen die in dit proefschrift worden beschreven hebben belangrijke implicaties voor de waardering van aandelen. Hoofdstuk 2 laat zien dat het bestaan van variatie in risico tot een verbetering leidt van de prestaties van het Fama en French (1993) driefactor model bij het verklaren van tijdsvariatie in rendementen op portefeuilles die gesorteerd zijn op de marktwaarde en de verhouding tussen de boek- en marktwaarde van een bedrijf. Voor sommige portefeuilles blijven waarderingsfouten echter groot en voorspelbaar. Bovendien kan het conditionele driefactor model de verklarende kracht van momentum variabelen voor de cross-sectie van rendementen niet wegnemen. In hoofdstuk 3 ligt de focus op individuele aandelen, waarvan de bèta's sterkere tijdsvariatie vertonen dan die van portefeuilles omdat de meeste veranderingen in het risico van individuele aandelen elkaar opheffen op portefeuilleniveau. De belangrijkste reden voor het samenvoegen van aandelen in portefeuilles is het verminderen van meetfouten in de schattingen van bèta's. De cruciale aanname van deze methode is dat alle aandelen die gegroepeerd worden in een bepaalde portefeuille dezelfde risicokenmerken hebben.

Hoofdstuk 3 levert sterk bewijs van cross-sectionele variatie in bedrijfsspecifieke bèta's binnen elk van de 25 portefeuilles gevormd op basis van marktwaarde en de verhouding tussen boek- en marktwaarde. Deze bevinding heeft belangrijke consequenties voor cross-sectionele toetsen van waarderingsmodellen. Door het groeperen van aandelen in portefeuilles wordt namelijk de cross-sectionele spreiding in bèta kleiner, waardoor informatie verloren gaat. Het gebruik van individuele aandelen in plaats van portefeuilles leidt daarom tot preciezere schattingen van risicopremies. In hoofdstuk 3 wordt een Bayesiaans panel model opgezet dat de informatie in de cross-sectie van aandelen gebruikt om de bèta's van individuele aandelen nauwkeuriger te schatten. Hiërarchische prior verdelingen worden gespecificeerd die een gemeenschappelijke structuur opleggen op de parameters van het model maar toch heterogeniteit op individueel niveau toestaan. De betere specificatie en preciezere schatting van tijdsvariërende bèta's leidt tot een grote toename in het waarderingsvermogen van het conditionele CAPM. De schatting van de marktrisicopremie is positief en significant en wordt niet beïnvloed door het

opnemen van bedrijfskarakteristieken in de cross-sectionele regressies. De bèta's uit het panel model worden ook gebruikt om de covariantiematrix van aandelenrendementen te voorspellen. Portefeuilles met minimale variantie gebaseerd op deze covariantiematrix zijn superieur in termen van *out-of-sample* standaarddeviatie aan portefeuilles die geconstrueerd zijn op basis van andere methoden.

Hoofdstuk 4 analyseert de waardering van plotselinge veranderingen in lange en korte termijn marktrisico's, individuele risico's en correlaties. De geïmpliceerde markt variantie berekend op basis van index opties overschrijdt systematisch de gerealiseerde markt variantie berekend met behulp van hoogfrequente rendementen. Dit wijst op de aanwezigheid van een variantierisicopremie in index opties. De decompositie van markt variantie in individuele varianties en correlaties laat zien dat variantierisico geprijsd is in index opties omdat correlatierisico geprijsd is. Daarnaast blijkt dat de markt variantierisicopremie een zeer goede voorspeller is van maand- en kwartaalrendementen op aandelenmarkten. Nadere analyses wijzen uit dat deze voorspellende kracht volledig wordt gedreven door de premie voor correlatierisico. Hoofdstuk 4 laat verder zien dat de prijs van marktvolatiliteitsrisico negatief is in de cross-sectie van aandelenrendementen, omdat beleggers een voorkeur hebben voor aandelen die het goed doen bij een onverwachte stijging in algemene onzekerheid. Lange termijn marktvolatiliteitsrisico is geprijsd omdat innovaties in zowel de gemiddelde lange termijn idiosyncratische volatiliteit als de gemiddelde lange termijn correlatie geprijsd zijn. Innovaties in korte termijn marktvolatiliteit zijn daarentegen alleen geprijsd omdat korte termijn correlatierisico geprijsd is. De variantie- en correlatiefactoren hebben verklarende kracht voor aandelenrendementen bovenop de verklarende kracht van traditionele risicofactoren, wat suggereert dat ze andere bronnen van systematisch risico oppikken.

Hoofdstuk 5 focust op de prestaties van particuliere beleggers in optiemarkten en onderzoekt of deze handelaren zich rationeel gedragen en de risico- en rendementskarakteristieken van opties begrijpen. Het onderzoek laat zien dat particuliere beleggers die opties verhandelen significant slechtere resultaten behalen dan beleggers die alleen in aandelen handelen. De verliezen van optiehandelaren zijn het gevolg van hoge transactiekosten en overreactie op grote koersdalingen in aandelenmarkten. Een gedetailleerde analyse van de transacties en bezittingen van optiebeleggers wijst uit dat de meerderheid van de beleggers *out-of-the-money* opties verhandelt zonder het onderliggende aandeel in bezit te hebben. Dit suggereert dat de meeste particuliere beleggers opties gebruiken om hun risico's te verhogen in plaats van te verlagen. Ondersteunend bewijs dat particulieren opties vooral verhandelen om te gokken is verkregen door het handelen in opties te relateren aan een aantal kenmerken van beleggers. Zo blijkt dat alleenstaande mannen met een laag inkomen en weinig beleggingservaring eerder geneigd zijn om opties te verhandelen. Het is bekend dat deze groep mensen ook een grotere neiging heeft om deel te nemen aan loterijen. Hoewel de prestaties van de gemiddelde optiebelegger teleurstellend zijn, ontdekt hoofdstuk 5 een klein aantal handelaren die erin slagen om consistent alle andere beleggers te verslaan. De groep succesvolle optiebeleggers handelt minder, is meer ervaren en bestaat uit meer vrouwen dan de groep handelaren die de grootste verliezen lijdt.

De inzichten die voortkomen uit het onderzoek in deze dissertatie zijn relevant voor academici, beleggers, managers en beleidsmakers. Allereerst zijn de resultaten van belang voor risicomangers in financiële instellingen omdat het risico van hun portefeuilles doorgaans binnen voorgeschreven limieten dient te blijven. Het bewijs van tijdsvariërende risico's in de hoofdstukken 2, 3 en 4 impliceert dat managers voortdurend de risico's van hun portefeuilles moeten monitoren en hun bezittingen aan moeten passen indien nodig.

Ten tweede zijn precieze schattingen van risico belangrijk voor het beoordelen van beleggingsprestaties. Als het waarderingsmodel dat gebruikt wordt om prestaties te meten geen rekening houdt met tijdsvariërende risico's of met belangrijke risicofactoren zoals de variantie- en correlatiefactoren uit hoofdstuk 4, dan kunnen fondsbeheerders uitzonderlijke rendementen behalen zonder daadwerkelijk over beleggingsvaardigheden te beschikken.

Ten derde vereist de implementatie van de portefeuilletheorie schattingen van verwachte varianties en covarianties, die lastig te meten zijn wanneer het beleggingsuniversum groot is. Het Bayesiaanse panel model in hoofdstuk 3 kan gebruikt worden om de precisie van de schattingen van individuele bèta's te vergroten en daardoor de accuratesse van voorspellingen van de covariantiematrix.

Ten vierde is een nauwkeurige schatting van de bèta van een bedrijf een belangrijke vereiste voor succesvolle investeringsbeslissingen omdat de verwachte toekomstige kasstromen van een project verdisconteerd moeten worden tegen de voor risico gecorrigeerde rentevoet van de onderneming. Inaccurate schattingen van bedrijfsspecifieke bèta's leiden tot gebrekkige schattingen van kapitaalkosten en daardoor tot onjuiste waarderingen van projecten.

Tot slot heeft de bevinding in hoofdstuk 5 dat cliënten van een online broker die in opties handelen zware verliezen lijden implicaties voor beleidsmakers. Dit is vooral van belang omdat blijkt dat de beleggers die het meeste handelen en het meeste verliezen het laagste inkomen hebben. Andere studies laten zien dat armere en lager opgeleide huishoudens meer beleggingsfouten maken, zoals het niet of nauwelijks beleggen in aandelenmarkten en het onvoldoende diversificeren van hun portefeuilles. Om mensen voor deze fouten te behoeden, moeten overheden trachten om de financiële kennis van de bevolking te vergroten door het opzetten van opleidingsprogramma's.

Deze dissertatie werpt ook een aantal nieuwe vragen op voor vervolgonderzoek. Allereerst moet er meer werk worden gedaan om de onderliggende economische bronnen aan het licht te brengen van de lange en korte termijn variantie- en correlatierisico's die in hoofdstuk 4 worden besproken. Daarnaast dienen er structurele modellen te worden ontwikkeld om de wig tussen geïmpliceerde en gerealiseerde varianties en correlaties te verklaren. Een andere mogelijke uitbreiding van de analyses in hoofdstuk 3 en 4 is het modelleren van lange en korte termijn fluctuaties in bèta's. Ook is het nog niet duidelijk of het irrationele gedrag van particuliere beleggers beschreven in hoofdstuk 5 een verstorende invloed heeft op optieprijsen. Omdat de resultaten in dat hoofdstuk suggereren dat veel beleggers de risico- en rendementskenmerken van opties onvoldoende begrijpen, rijst bovendien de vraag of beleggers van hun fouten leren en of scholing hun prestaties kan verbeteren.

Curriculum Vitae

Mathijs Cosemans was born on June 23, 1982 in Born, the Netherlands. He attended high school between 1994 and 2000 at the Trevianum Gymnasium in Sittard. Subsequently, he studied Economics and Econometrics at Maastricht University. In April 2005 he received his Master's degree in Financial Economics with distinction and in September 2005 he obtained his Master's degree in Econometrics and Operations Research. He received the best Master's thesis award in Economics from Maastricht University in 2005.

In September 2005 Mathijs joined the Department of Finance at Maastricht University as a Ph.D. candidate. The results of his work on risk and return dynamics in stock markets are presented in this dissertation. He performed part of his research while attending Harvard University as a visiting scholar. He was also a teaching assistant for several courses in investments and financial econometrics.

His work has been presented at various universities and international conferences, including Harvard University, Yale University, the Stockholm School of Economics, the University of Amsterdam, and the Annual Meetings of the American Finance Association, Western Finance Association, European Finance Association, Financial Management Association, Centre for Economic Policy Research, Econometric Society, and Society for Financial Econometrics. A research proposal for one of the chapters in his dissertation was awarded a grant from Inquire Europe. His work has been published or is forthcoming in international refereed academic journals. In December 2009 Mathijs joined the Finance Group at the University of Amsterdam as a postdoctoral research fellow.