Essays on Hedge Fund Illiquidity, Return Predictability, and Time-Varying Risk Exposure



Mathias S. Kruttli
Lincoln College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy in Economics

Hilary 2015

Acknowledgements

I would like to thank my supervisors Tarun Ramadorai and Kevin Sheppard. I have profited tremendously from their guidance and encouragement throughout my studies. Thanks to their continuous feedback, I have had a very smooth transition from the MPhil to the DPhil and constantly improved my research and teaching skills. Further, I highly appreciate their support regarding the next steps of my academic career. This thesis and what follows would not have been possible without them.

I also thank my co-author Andrew Patton. Working with Andrew and Tarun on our paper was a fantastic experience. This experience has not only taught me how to conduct research in empirical finance at a very high standard, but has also provided me with an invaluable exposure to the international academic finance community.

I have had the pleasure to meet numerous people who contributed to my time as a doctoral student through their academic expertise, company, and/or tennis skills. Many thanks to: Mahmood Al Saleh, Haarpret Bedi, Ton van den Bremer, Martin Fluder, Daniel Gassmann, Irem Guceri, Alex Hadcock, Moritz Hafner, Jan Jozwik, Sebastian Koenigs, Kasper Lund-Jensen, Zsolt Mester, Diaa Noureldin, Cavit Pakel, Anders Rahbek, Max Roser, Michael Streatfield, Andreas Tischbirek, Rasmus Varneskov, and James Wolter. In particular, I would like to thank Sumudu Watugala for making my time at the University of Oxford truly exceptional.

I gratefully acknowledge financial support from the David Walton Memorial Fund, Lincoln College, the Oxford-Man Institute of Quantitative Finance (OMI), the Swiss National Science Foundation, and the University of Oxford Economics Department. I am particularly thankful to the OMI for providing a terrific research environment.

Finally, I am deeply grateful to my parents Markus and Regina and my sister Vera for their support throughout my life.

Abstract

This thesis consists of three papers that make independent contributions to the field of financial economics. As such, the papers, Chapter 2, 3, and 4, can be read independently of each other.

In Chapter 2, we construct a simple measure of the aggregate illiquidity of hedge fund portfolios, and show that it has strong in- and out-of-sample forecasting power for 72 portfolios of international equities, U.S. corporate bonds, and currencies, over the 1994 to 2011 period. The forecasting ability of hedge fund illiquidity for asset returns is, in most cases, greater than, and provides independent information relative to, well-known predictive variables for each of these asset classes. We construct a simple equilibrium model to rationalise our findings and empirically verify auxiliary predictions of the model.

In Chapter 3, I analyse the risk-shifting of hedge funds. Since the information on hedge fund holdings is very restricted, researchers have used the variance of returns as a proxy for risk. I propose a new method for measuring the time-varying variance. I use this method to investigate whether equity long-short hedge funds engage in risk-shifting driven by their past performance relative to their peers. I find that hedge funds which have strongly underperformed or outperformed their peers in recent months increase their exposure to the core strategy, i.e. the equity long-short strategy, and to non-core strategies. The risk shifting is mitigated for hedge funds with long redemption periods.

Chapter 4 contributes to the equity premium prediction literature. I improve the forecast performance of typical single variable predictive regressions used in the equity premium prediction literature through Bayesian priors derived from consumption-based asset pricing models. To implement these model-based priors, I develop a Bayesian procedure which is rooted in the macroeconometrics literature. I find that the model-based priors can increase the explanatory power, measured by the out-of-sample R^2, of the single variable predictive regressions by several percentage points.

Contents

1 Introduction **1**

2 The Impact of Hedge Funds on Asset Markets **4**

2.1 Introduction 5

2.2 Data Description 8

2.2.1 Asset return data 8

2.2.2 An index of hedge fund illiquidity 10

2.3 Empirical Results on Asset Returns and Hedge Fund Illiquidity . . . 16

2.3.1 Does hedge fund illiquidity predict asset returns? 16

2.3.2 The out-of-sample predictive power of hedge fund illiquidity . 23

2.4 A Model of Hedge Fund Illiquidity and Asset Return Determination 29

2.4.1 The timeline of the model 29

2.4.2 The hedge fund's objective function 30

2.4.3 Market clearing and the equilibrium price 31

2.4.4 Model-implied hedge fund illiquidity and asset returns 32

2.4.5 Empirical predictions of the model 37

2.5 Robustness Checks 41

2.6 Conclusion 44

Appendices **45**

2.A Extension of Clark and West (2006) 45

2.A.1 Properties of the extension 45

2.A.2 Monte Carlo simulation study 49

2.B Hedge Fund Style Indices vs. an Aggregate Index 53

2.C Numerical Procedure to Simulate Model 60

2.D Additional Tables 63

3 On the Time-Varying Risk of Hedge Fund Portfolios and Risk-Shifting 66
3.1 Introduction . . . 67
3.2 Data . . . 69
3.2.1 Hedge fund returns . . . 69
3.2.2 Equity long-short strategy returns . . . 70
3.3 Methodology . . . 72
3.3.1 Rolling window estimation . . . 73
3.3.2 Dynamic conditional estimation . . . 74
3.3.2.1 Estimation of the conditional variances . . . 75
3.3.2.2 Estimation of the conditional correlation . . . 77
3.3.2.3 Empirical tests to detect risk-shifting . . . 78
3.4 Results . . . 80
3.4.1 Results for the rolling window estimation . . . 80
3.4.2 Results for the dynamic conditional estimation . . . 81
3.5 Robustness Checks . . . 92
3.6 Conclusion . . . 94

Appendices 97
3.A The Panel AVGARCH(1,1) Model . . . 97

4 Model-Based Priors for Predicting the Equity Premium 98
4.1 Introduction . . . 99
4.2 Bayesian Motivation . . . 101
4.3 Econometric Methodology . . . 103
4.3.1 Empirical procedure . . . 103
4.3.2 Model-based prior . . . 105
4.4 Data . . . 107
4.4.1 Empirical data . . . 108
4.4.2 Simulated data . . . 108
4.4.2.1 Monte Carlo simulation . . . 109
4.4.2.2 Habit Formation model . . . 109
4.4.2.3 Prospect Theory model . . . 110
4.4.2.4 Long Run Risk model . . . 110
4.4.2.5 Habit Formation Term Structure model . . . 111
4.5 Results . . . 112
4.5.1 Habit Formation model-based priors . . . 112
4.5.2 Prospect Theory model-based priors . . . 114

4.5.3 Long Run Risk model-based priors . . . 114
4.5.4 Habit Formation Term Structure model-based priors . . . 119
4.5.5 Sensitivity of the model-based prior's OOS performance to λ . 120
4.5.6 Is the model-based prior truly OOS? . . . 121
4.5.7 The model-based prior versus the Minnesota prior . . . 124
4.6 How the Model-Based Prior Improves Predictive Accuracy . . . 127
4.7 Conclusion . . . 133

Appendices **135**
4.A Model-Based Priors . . . 135
4.A.1 The prior and posterior distributions . . . 135
4.A.2 Normalisation factors . . . 138
4.A.3 Data densities . . . 138
4.B Asset Pricing Models . . . 139
4.B.1 By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior . . . 139
4.B.2 A Consumption-Based Model of the Term Structure of Interest Rates . . . 141
4.B.3 Prospect Theory and Asset Prices . . . 143
4.B.4 Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles . . . 145
4.C Additional Tables . . . 148

5 Conclusion **149**

Bibliography **152**

List of Figures

2.1 Hedge Fund Illiquidity Index from January 1995 to December 2011 . 14
2.2 In-Sample Single Variable Predictability at Different Horizons 24
2.3 In-Sample Multiple Variables Predictability at Different Horizons . . 24
2.4 Model Implied Return Reversal 36
2.5 Test Statistic Distribution of the Clark-West Extension 52

3.1 Average β_t Estimates over Time 83
3.2 Average $\tilde{\sigma}_t$ Estimates over Time 83

4.1 Sensitivity of R^2_{OS} to λ 122
4.2 Parameter Estimates over the Postwar Sample 129
4.3 Cumulative Sum of Squared Errors over the Postwar Sample 132

List of Tables

2.1 Summary Statistics of Asset Returns 11
2.2 Summary Statistics of Hedge Fund Illiquidity Index 15
2.3 In-Sample Predictive Performance 19
2.4 Out-of-Sample Forecast Performance 27
2.5 Model Simulation 34
2.6 Panel Estimation Conditioning on Sign of Lagged Return 38
2.7 Panel Estimation with Illiquidity Dummy Variables 40
2.8 Robustness Checks 42
2.9 Rejection Probabilities for the Clark-West Extension 51
2.10 In-Sample Predictive Performance for U.S. Corporate Bonds 63

3.1 Summary Statistics 71
3.2 Panel Model With Rolling Window Estimates 82
3.3 Volatility Estimates 85
3.4 Panel Model With Dynamic Conditional Estimates 89
3.5 Panel Model With Dynamic Conditional Estimates and Redemption Dummy 90
3.6 Panel Model With Dynamic Conditional Estimates and HWM Benchmark 91
3.7 Panel Model With Dynamic Conditional Estimates and Quarterly Dummy 93
3.8 Robustness Checks 95

4.1 Results for the Habit Formation Model-Based Priors 115
4.2 Results for the Prospect Theory Model-Based Priors 116
4.3 Results for the Long Run Risk Model-Based Priors 117
4.4 Results for the Habit Formation Term Structure Model-Based Priors 118
4.5 Results for the Post Calibration Period 125
4.6 Comparison of the Model-Based Prior and the Minnesota Prior . . . 128
4.7 Parameter Estimates for the Postwar Sample 130

4.8 Comparison of the Model-Based Prior and the Campbell and Thompson Restrictions 148

Chapter 1

Introduction

The predictability of asset returns is a highly debated topic among researchers in asset pricing and financial econometrics. The latest manifestation of this debate has been the award of the The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2013 to researchers whose answers to the question, whether returns are predictable or not, appear to differ. An argument against the predictability of returns is that "smart" investors trade against any mispricing in financial markets and make predictability patterns vanish quickly. A counterargument by proponents of predictable returns is that limits of arbitrage can prevent smart investors from executing rapid and profitable trading to bring prices in line with fundamentals.

Chapter 2 of this thesis contributes to this ongoing debate. We investigate if the portfolio illiquidity of hedge funds, which are widely considered to be smart investors, can help us to forecast asset returns. We estimate an aggregate index of hedge fund illiquidity through the autocorrelation of hedge fund returns reported at a monthly frequency. The theory predicts that high autocorrelation implies high illiquidity due to return smoothing, non-synchronous trading, and extrapolation of past returns. We find that our index of hedge fund illiquidity indeed has strong forecasting power for the returns of a range of asset classes: high hedge fund illiquidity predicts high returns. We explain this finding with a theoretical model in which a hedge fund acts as a market maker for (and trades against the mispricing of) a risky asset. If the hedge fund has an illiquid portfolio, the market maker activities

are constrained: the hedge fund will be reluctant to buy the risky asset but eager to sell it. The consequent price drop makes the return reversal – also interpreted as the compensation to the market maker – greater when the hedge fund buys the risky asset and smaller (in absolute values) when the hedge fund sells.

The third chapter also focuses on hedge funds. Hedge fund data can pose challenges due to short time series: hedge fund data are not available before January 1994, and low reporting frequency. One challenge is the estimation of the time-varying risk exposure. Since there is little transparency in the asset allocation of institutional investors, the variance of the portfolio returns is often used as a proxy for portfolio risk in the hedge fund and mutual fund literature. For mutual funds, returns are often reported daily, which makes the realised variance a viable estimator for the monthly variance – the realised variance is computed with an average of approximately 22 data points for each month. However, for hedge funds, only monthly return data are available. How can we estimate a time-varying monthly variance with monthly data? Estimating a realised variance with rolling windows is one solution, but a solution with serious drawbacks. A major drawback is that the window length is often chosen on an ad-hoc basis and estimates can vary greatly across different window lengths. I propose a new method for estimating the time-varying risk exposure of hedge funds. This method combines recent advances in the conditional variance and correlation estimation literature. The estimates of time-varying risk are then used to analyse whether hedge funds engage in risk-shifting when underperforming or outperforming their peers in year-to-date returns. In line with the existing literature on hedge fund and mutual fund tournaments, I find that equity long-short funds take more risk the farther they lag behind their peers. However, my measure of time-varying risk exposure also reveals evidence of a similar mechanism for outperforming hedge funds: hedge funds which have performed well relative to their peers tend to increase risk over the subsequent months.

In the fourth chapter, I revisit the topic of return predictability. In the equity

premium prediction literature, predictor variables, for example the dividend-price ratio or the dividend yield, are used to predict the difference in returns on the S&P 500 index and the risk-free rate. A problem which arises in out-of-sample exercises is the overfitting of parameter estimates – the estimates pick up noise which leads to imprecise forecasts of returns not included in the estimation sample. To overcome this problem, researchers have imposed restrictions on the parameter estimates. Chapter 4 contributes to this literature by adapting a technique from macroeconometrics which uses Bayesian priors derived from dynamic stochastic general equilibrium models and imposes them on vector autoregressions to reduce the parameter overfitting. Correspondingly, my adaptation derives priors from consumption-based asset pricing models, which are then imposed on the predictive regressions to forecast the equity premium. These priors, derived from models whose purpose is to provide an explanation for the equity premium puzzle, lead to an increase in the out-of-sample R^2 of several percentage points. Further, this methodology allows a comparison of the consumption-based asset pricing models across a new dimension, namely the effectiveness of their priors for the prediction of returns.

Chapter 2

The Impact of Hedge Funds on Asset Markets

Abstract[1]

While there has been enormous interest in hedge funds among academics, prospective and current investors, and policymakers, rigorous empirical evidence of their impact on asset markets has been difficult to find. We construct a simple measure of the aggregate illiquidity of hedge fund portfolios, and show that it has strong in- and out-of-sample forecasting power for 72 portfolios of international equities, U.S. corporate bonds, and currencies, over the 1994 to 2011 period. The forecasting ability of hedge fund illiquidity for asset returns is, in most cases, greater than, and provides independent information relative to, well-known predictive variables for each of these asset classes. We construct a simple equilibrium model to rationalise our findings and empirically verify auxiliary predictions of the model.

[1]This chapter is entirely based on work co-authored with Andrew Patton and Tarun Ramadorai.

2.1 Introduction

Hedge funds have sparked enormous interest among a number of different constituencies. Wealthy individual investors and pools of institutional capital are interested in them as a potential source of high returns with the promise of low risk. Academics have been deeply interested in whether these high returns are persistent and whether there are important risk exposures underlying these returns, in addition to studying hedge fund fees, capital accumulation, and disclosure policies. Regulators and policymakers have been wary about the industry ever since the Long-Term Capital Management (LTCM) crisis of 1998. The periodic collapses of major hedge funds such as Amaranth and Madoff have perpetuated these concerns, because of the perception that such collapses could have an impact on the underlying asset markets in which these funds are invested. These concerns are understandable: while the global hedge fund industry has only around U.S.$ 1.5 trillion of assets under management (AUM), hedge funds' substantial leverage and the high levels of trading volume that they generate in underlying asset markets means that their impact may well be disproportionately large.

Despite this high level of interest and concern, there has been a paucity of compelling empirical evidence that connects the activity of hedge funds to returns in the underlying asset markets. We attempt to fill this gap by providing evidence that hedge funds have an important, measurable, impact on expected returns. We do so by constructing a simple measure of their aggregate ability to provide liquidity to asset markets, and show that this measure has strong predictive power for a wide range of assets spanning three broad categories: international equities, U.S. corporate bonds, and currencies.

Hedge funds are often characterised as "arbitrageurs" in financial markets. Canonically, this term refers to the generation of a riskless return by the simultaneous trade of under- and over-valued securities. However, the term is often abused when applied to hedge funds, and generally describes the exploitation of profitable, albeit

risky investment opportunities, with an emphasis on managing risk to the greatest possible extent. One particularly important source of such investment opportunities for hedge funds has been the provision of liquidity to asset markets (see, for example, Aragon (2007), Sadka (2011), and Jylha, Rinne, and Suominen (2013)). As a result of exploiting these opportunities, hedge funds are exposed to illiquidity risk, and impose investment restrictions on their outside investors to safeguard against the withdrawal of hedge fund capital over short horizons when the illiquidity-driven price pressure is probably most acute. The extent of their exposure can be measured: Getmansky, Lo, and Makarov (2004) cleverly identify that hedge funds holding the most illiquid investments are likely to exhibit persistent returns, and propose that the extent of the illiquidity for any given fund at a point in time can be ascertained by the autocorrelation of its returns.

We take this insight as the starting point for our analysis. We measure the extent of return autocorrelation for each fund in our comprehensive universe of roughly 30,000 hedge funds, in each month over the period from 1994 to 2011. We then aggregate this measure by simply averaging the fund-specific first-order autocorrelation estimates each month. We find that the resulting measure of the aggregate illiquidity of hedge funds, which we dub "ρ", has strong, consistent predictive power for 72 portfolios of assets spanning three major asset classes. In-sample, our measure is significant for 21 out of 21 international equity indices, 31 out of 42 U.S. corporate bond portfolios spanning the ratings and maturity spectrum, and 6 out of 9 developed-world currencies that we consider. This is strong, broad support for the importance of our new measure.

This predictive power of ρ is not just in-sample; we also find evidence that the measure has appreciable out-of-sample (OOS) forecasting power. Out-of-sample, we can reject the historical mean return model with our predictor for 20 out of 21 international equity indices, 28 out of 42 U.S. corporate bond portfolios, and 3 out of 9 currencies. In both the in-sample and out-of-sample predictive regressions,

our proposed measure generally outperforms, and always contributes incremental explanatory power relative to a range of competitor variables, which are known from the relevant asset class-specific literature to be useful forecasters of asset returns.

To rationalise our findings, we build a simple equilibrium model, which incorporates liquidity constraints into the limits of arbitrage framework of Gromb and Vayanos (2010). In our model, the hedge fund begins with an endowment of illiquid assets and cash, and makes returns by providing liquidity to absorb buying and selling pressure from noise traders. The hedge fund thus acts as a quasi-market-maker for the risky asset. However, the liquidity provision capacity of the hedge fund is limited by the threat of its outside investors withdrawing their funds, which forces it to hold a sufficient quantity of cash to satisfy these redemptions. This means that an illiquid hedge fund is (relatively) reluctant to purchase the risky asset and (relatively) eager to sell it. We evaluate the comparative statics in the model by varying the ratio of the illiquid endowment and cash held by the hedge fund and find that this has an impact on the expected returns – the higher the measure of illiquidity of the hedge fund's portfolio, the higher the expected asset returns.

Following the logic of our model, the return predictability of ρ should be stronger for illiquid assets: an illiquid hedge fund is even more reluctant to purchase the risky asset and is inclined to sell it, when the asset itself is illiquid. We verify this in empirical tests. We find the highest levels of predictability for low market captilisation and low turnover equity markets, high-yield and long maturity corporate bonds, and currencies with high short term interest rates and bid-ask spreads.

The literature on hedge funds is fast-growing – many authors, such as Fung and Hsieh (1997, 2001, 2004), Agarwal and Naik (2004), Mamaysky, Spiegel, and Zhang (2007), Bollen and Whaley (2009), and Patton and Ramadorai (2013), have documented that hedge funds are significantly exposed to systematic risk, proxied by indices of equity, bond, and options returns. Rigorous empirical evidence on the reverse direction, namely the impact of hedge funds on asset markets, has been less

well-documented. For example, Kang, Kondor, and Sadka (2012) provide evidence suggesting that hedge funds affect idiosyncratic risk in equity markets. This chapter provides wide-ranging evidence about hedge funds' impact on a variety of asset markets using a relatively simply constructed variable, which we view as a useful contribution to this emerging literature.

This chapter also adds to the literature on the impacts of liquidity and liquidity risk in asset markets by contributing an important new variable: aggregate hedge fund illiquidity. We view our newly introduced variable as a complement to a number of asset-class-specific liquidity measures (see, for example, Pastor and Stambaugh (2003), Acharya and Pedersen (2005), Khandani and Lo (2011), and Bongaerts, de Jong, and Driessen (2012)) that have been utilised extensively in the literature.

The remainder of the chapter is organised as follows. Section 2.2 describes the construction of our hedge fund illiquidity index as well as the asset return data employed in our study. Section 2.3 presents the in-sample and out-of-sample forecasting results using the illiquidity index. In Section 2.4, we outline our model and empirically test additional model predictions. Section 2.5 considers the robustness of our results to a number of different variations, and the final section concludes.

2.2 Data Description

2.2.1 Asset return data

Our empirical analysis covers three major asset classes: international equities, U.S. corporate bonds, and currencies. Within each asset class, we study a broad range of individual assets. As our hedge fund data are only available monthly, and our total time series length is under 20 years, we focus on monthly asset return predictability rather than higher- or lower-frequency returns. For each of the asset classes, we compare the predictive performance of our hedge fund illiquidity measure with that of a range of benchmarks from the extant literature. We describe this set of

"competing" predictor variables utilised for each asset class below.

We employ 21 international equity index returns in our analysis, namely Austria, Australia, Belgium, Canada, Denmark, Finland, France, Germany, Hong Kong, Ireland, Italy, Japan, Netherlands, New Zealand, Norway, Singapore, Spain, Sweden, Switzerland, U.K., and U.S. We source our data on returns from Kenneth French's website. We compute log excess returns on each of these indices over the sample period 1995 to 2011. As competitor predictor variables, we employ the dividend yield for each equity index from Kenneth French's website, as well as innovations in the VIX index, computed using an AR(2) model.[2]

Our dataset on U.S. corporate bonds comprises 42 indices from the Bank of America-Merrill Lynch, of which 24 are investment-grade portfolios and 18 are high-yield portfolios, with maturities ranging from one year to over fifteen years. Our corporate bond data begin in 1997 and end in 2011. Following Bongaerts et al. (2012), we compare the return prediction performance of our hedge fund illiquidity index with that of the Pastor-Stambaugh traded liquidity factor acquired from CRSP, innovations to the VIX, estimated as described above, and market-capitalisation weighted excess returns on the S&P 500 index, also from CRSP.

Our data on currencies comprise nine currency rates against the U.S. Dollar: Australian Dollar, Canadian Dollar, Euro, Japanese Yen, New Zealand Dollar, Norwegian Krone, Swedish Krona, Swiss Franc, and British Pound. The data are sourced from Bloomberg and cover the period 1995 to 2011 (we use the Deutsche Mark prior to the introduction of the Euro). Our monthly log currency returns are measured as U.S. Dollars per unit of foreign currency.[3] Given the findings of Meese and Rogoff (1983), we expect that these returns will be extremely difficult to predict. The competitor variables for predicting currency returns include the inflation differential

[2]Goyal and Welch (2008) consider a wider array of predictor variables than those listed. However our focus on international equity returns, rather than just U.S. returns, means that the number of available competitors is restricted by data limitations.

[3]We also considered excess currency returns (over the interest rate differential) and found that the results were very similar.

from the OECD database and the one month Libor interest rate differential from Bloomberg.

Summary statistics for all three asset classes can be found in Table 2.1. Given the large number of individual bond return series (42 in total), we group them into investment grade and high yield bonds as well as separate them by maturity.

2.2.2 An index of hedge fund illiquidity

To compute our hedge fund illiquidity index, we employ monthly hedge fund returns over the period from January 1994 to December 2011, consolidated from data in the TASS, HFR, CISDM, Morningstar, and BarclayHedge databases. These data represent the most comprehensive set of hedge fund data available from public sources, and comprise a total of 29,496 individual hedge funds, including all births and deaths of funds over the period. Details of the process followed to consolidate these data can be found in Patton, Ramadorai, and Streatfield (2013). Reported returns are net of management and incentive fees.

Getmansky et al. (2004) and Lo (2007) propose using autocorrelation in a hedge fund's returns as a proxy for the illiquidity of its asset holdings. Their rationale is that hedge fund managers mark their portfolios to market at the end of each month, and may be forced to use models to estimate portfolio values owing to the illiquidity of the assets in the portfolio. One such simple model is the linear extrapolation of returns in the portfolio, which would lead to reported hedge fund returns exhibiting a positive autocorrelation. Another factor generating spurious positive autocorrelation is holding both liquid and illiquid assets in a portfolio, generating non-synchroneity of returns (see, for example, Scholes and Williams (1977) and Dimson (1979)). Thus autocorrelation in hedge fund returns provides a (noisy) measure of the degree of illiquidity of the hedge fund's holdings – higher levels of autocorrelation are associated with higher levels of illiquidity.

Getmansky et al. (2004) also consider the case that hedge fund managers engage

Table 2.1: Summary Statistics of Asset Returns

Reported are the average, the standard deviation, the maximum, and the minimum, for monthly log excess returns in % of international equities and U.S. corporate bonds (values are computed across the respective portfolios for U.S. corporate bonds). For currencies, the statistics are based on monthly returns with the U.S. Dollar as the base currency. For international equities and currencies, the time series starts in January 1995 and ends in December 2011. For U.S. corporate bonds, the time series starts in January 1997 and ends in December 2011.

Panel A: International Equities				
	Average	**Std. Dev.**	**Max**	**Min**
Australia	0.625	6.403	15.989	-32.581
Austria	0.339	7.058	17.756	-42.588
Belgium	0.356	6.020	15.331	-36.274
Canada	0.721	6.178	19.901	-31.179
Denmark	0.587	5.950	16.635	-28.925
Finland	0.575	9.239	26.508	-34.564
France	0.392	6.211	13.916	-24.795
Germany	0.316	6.679	19.782	-26.658
Hong Kong	0.401	7.560	27.622	-34.038
Ireland	0.325	6.356	16.202	-26.787
Italy	0.195	7.091	18.914	-27.153
Japan	-0.282	5.462	15.200	-14.452
Netherlands	0.365	6.551	15.452	-34.552
New Zealand	0.186	6.390	14.453	-21.282
Norway	0.592	7.891	17.681	-36.923
Singapore	0.209	7.792	25.642	-33.974
Spain	0.519	6.661	18.231	-26.032
Sweden	0.680	7.686	22.440	-32.318
Switzerland	0.495	5.044	12.130	-16.435
U.K.	0.379	4.766	13.208	-22.279
U.S.	0.448	4.622	10.356	-18.347
Panel B: U.S. Corporate Bonds				
Rating				
Inv. Grade (24 Portfolios)	0.346	1.703	15.146	-16.070
High Yield (18 Portfolios)	0.375	3.811	31.302	-36.268
Maturity				
1-3Y (7 Portfolios)	0.325	1.827	31.302	-22.596
3-5Y (7 Portfolios)	0.325	1.931	14.882	-24.184
5-7Y (7 Portfolios)	0.344	2.280	20.874	-25.359
7-10Y (7 Portfolios)	0.279	2.774	17.899	-30.674
10-15Y (7 Portfolios)	0.347	3.230	23.176	-36.268
15+Y (7 Portfolios)	0.532	3.596	24.717	-28.694
Panel C: Currencies				
Australia	0.135	3.653	9.899	-17.108
Canada	0.155	2.455	8.936	-13.009
Euro	0.019	3.017	9.609	-10.196
Japan	0.127	3.259	16.273	-9.698
New Zealand	0.096	3.640	12.507	-13.931
Norway	0.061	3.201	7.785	-13.739
Sweden	0.037	3.318	9.127	-11.621
Switzerland	0.163	3.244	12.368	-11.945
U.K.	-0.003	2.441	9.044	-10.215

in performance smoothing, whereby reported returns are a smoothed version of the true returns (see Bollen and Pool (2008, 2009) and Agarwal, Daniel, and Naik (2011) for related work). Intentional smoothing leads to the same features as the smoothing that arises from marking to model: reported returns have a positive autocorrelation even when the underlying true returns do not. If some managers engage in smoothing regardless of the market conditions, then this will work against us finding a significant result: the relationship between hedge fund portfolio illiquidity and asset returns will be muddied by individual fund illiquidity estimates that are unrelated to true illiquidity. However, if these managers engage in opportunistic smoothing – they smooth more when the assets in their portfolios are illiquid – then the degree of intentional smoothing will be correlated with the true portfolio illiquidity.

Our aggregate measure of hedge fund portfolio illiquidity, which we denote by ρ_t in month t, is computed using the following simple procedure. First, we compute the first-order autocorrelation of each hedge fund's returns using a rolling 12-month window. Having computed these return autocorrelations for all individual funds for which we have the prior 12 months of returns available in each month, we impose a lower bound of zero on any autocorrelation estimate which is estimated to be negative.[4] We then simply average these autocorrelations across all funds for each month in our sample t, yielding ρ_t. In our main analysis, we use an equally weighted average of these individual illiquidity estimates, and in our robustness checks, we show that using an AUM weighted measure leads to similar results.

Given the 12-month burn-in period for estimating rolling autocorrelations, our time series of the hedge fund illiquidity index begins in December 1994 and ends in December 2011. Figure 2.1 plots the AUM weighted and the equally weighted hedge fund illiquidity measures over time, and shows that the hedge fund illiquidity spikes during the LTCM crisis of 1998, the Quant Meltdown of August 2007 (see

[4]We impose the filter simply because the interpretation of positive autocorrelation estimates is conceptually far easier. In our robustness checks, we do not impose this filter, and find that our results remain strong. Estimation error can drive positive autocorrelation estimates below zero when the true values are above zero.

Khandani and Lo (2011)), as well as during the Great Recession. However, the recession following the NASDAQ crash at the turn of the millennium did not appear to affect the measure greatly. It is also evident from the plot that the equally weighted illiquidity index has far greater volatility than the value weighted index. This suggests that larger funds manage their liquidity more effectively. This may not be surprising given that funding liquidity pressure imposed by prime-brokerage relationships and capital inflows to funds may be less of an issue for larger, better-established hedge funds.

Summary statistics for the measure are reported in Table 2.2. The funds in the combined database come from a broad range of vendor-classified strategies, which we consolidate into nine main strategy groups: Directional Traders, Emerging Markets, Fixed Income, Funds-of-Funds, Macro, Managed Futures (CTA), Multi-Process, Relative Value, and Security Selection (see Patton et al. (2013) for detailed mappings between self-reported classifications and these broad categories).

The table reports summary statistics for hedge fund illiquidity indices computed for each individual style, followed by those computed across all funds, regardless of style. The table reveals cross-sectional variation in illiquidity that is consistent with intuition. For example, the Fixed Income style has a high mean level of illiquidity, while the CAT funds are the least illiquid. In our analysis, we focus on the illiquidity index based on all funds, rather than that derived from a specific style. This is partly motivated by a desire for simplicity, and partly by a desire for robustness to potential style drift and misclassification arising from the static and self-reported style classifications in our hedge fund databases. Appendix 2.B contains an analysis of how these features can lead to an index based on all funds that can deliver better results than one which relies upon noisy style classifications. In Section 2.5, we show the results obtained when employing illiquidity indices based on hedge funds in styles that are, in theory, more closely related to a given asset class, and find that they are similar.

Figure 2.1: Hedge Fund Illiquidity Index from January 1995 to December 2011

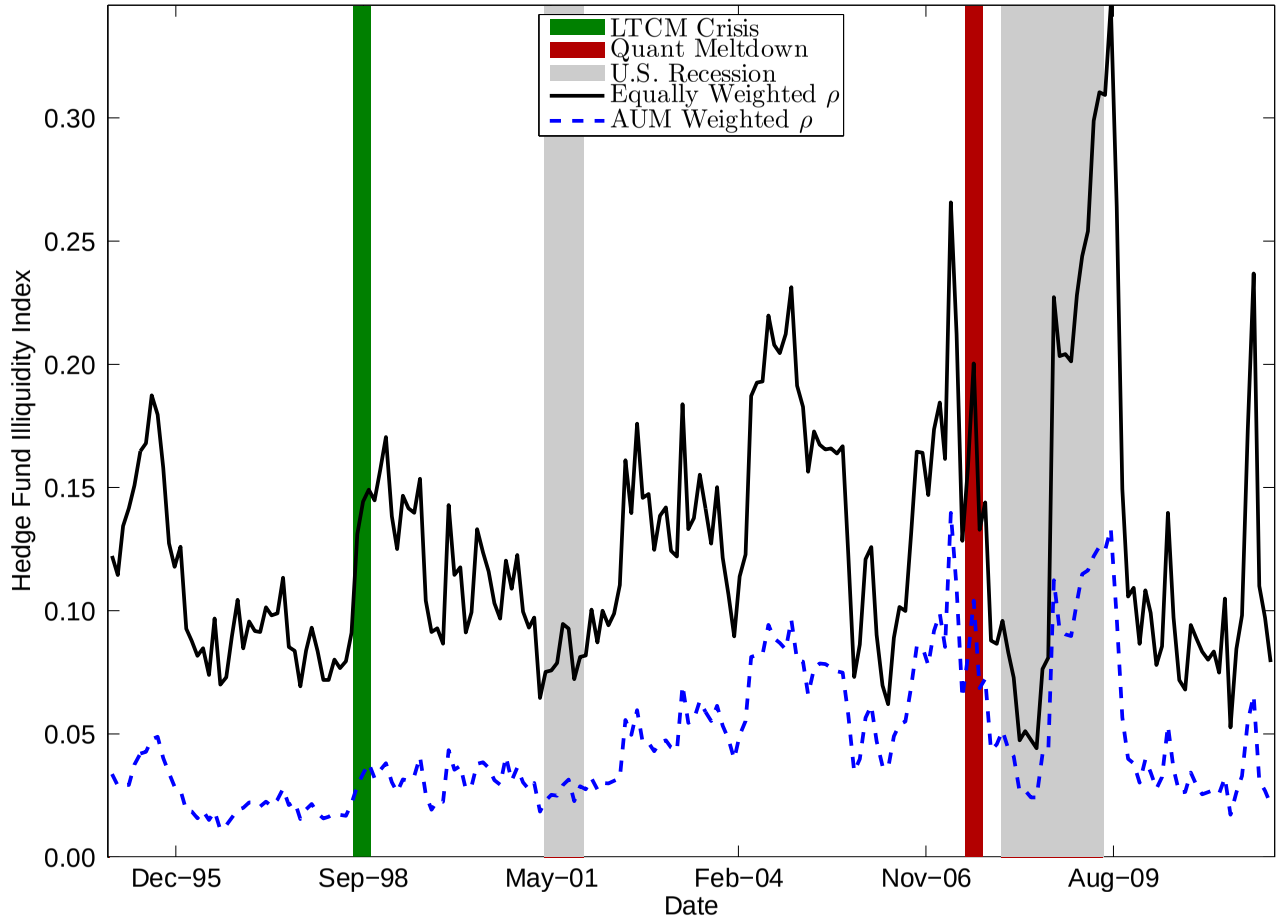


Table 2.2: Summary Statistics of Hedge Fund Illiquidity Index

For the illiquidity index of each investment style, we report the average, the standard deviation, and the correlation with the All hedge fund illiquidity index. The hedge fund illiquidity index for a specific hedge fund in month t is the autocorrelation of the hedge fund's monthly returns estimated over the last 12 months. The hedge fund illiquidity index shown in the table is an equally or AUM weighted average across all the hedge funds of a particular investment style. Negative autocorrelations are trimmed at zero. The time series is from January 1995 to December 2012.

	Equally Weighted			**AUM Weighted**		
Styles	**Average**	**Std. Dev.**	**Corr. with All**	**Average**	**Std. Dev.**	**Corr. with All**
CTA	0.086	0.035	0.718	0.022	0.013	0.586
Directional Traders	0.111	0.038	0.876	0.030	0.028	0.866
Emerging Markets	0.137	0.069	0.795	0.044	0.028	0.686
Fixed Income	0.174	0.051	0.786	0.056	0.025	0.746
Funds of Funds	0.154	0.089	0.958	0.062	0.053	0.972
Global Macro	0.094	0.034	0.775	0.028	0.018	0.821
Multi-Strategy	0.147	0.055	0.905	0.061	0.031	0.919
Relative Value	0.102	0.039	0.678	0.027	0.033	0.644
Security Selection	0.111	0.038	0.876	0.046	0.018	0.795
All	**0.127**	**0.053**	**1.000**	**0.046**	**0.027**	**1.000**

2.3 Empirical Results on Asset Returns and Hedge Fund Illiquidity

We now present our main analysis of the predictive power of the hedge fund illiquidity index for future asset returns. We study this relationship using a variety of simple models. Our first set of results are based on full-sample estimation of the relationship, and we conclude with an out-of-sample forecasting analysis.

2.3.1 Does hedge fund illiquidity predict asset returns?

Our first analysis of the relationship between hedge fund illiquidity and asset returns is based on a simple predictive regression, in which we use the hedge fund illiquidity measure described above, denoted ρ_t, to predict one-month-ahead returns $r_{i,t+1}$ on asset i:

$$r_{i,t+1} = \alpha_i + \gamma_i \rho_t + \epsilon_{i,t+1}, \text{ where } \epsilon_{i,t+1} \sim N(0, \sigma^2_{i,t+1}). \tag{2.1}$$

To account for heteroskedasticity and the serial correlation of the normally distributed error term, $\epsilon_{i,t+1}$, we use the standard errors proposed by Newey and West (1987). For comparison, we also estimate the same univariate model using various competitor variables, suggested by the asset-class-specific literature, in place of ρ_t. We also include hedge fund flows as a competitor, as they affect the liquidity of a hedge fund's portfolio, and Jylha et al. (2013) find that increases in hedge fund flows reduce the amount of short-term return reversal and volatility in the U.S. equity market. Moreover, our model in Section 2.4 discusses the different roles of flows and the starting endowment of a fund's illiquid assets, making flows a natural competitor for ρ_t. We compute these flows fund-by-fund using the standard approach (see, for example, Fung, Hsieh, Naik, and Ramadorai (2008)), winsorize them at the 1% and 99% percentile points, and value weight them to construct an aggregate flow index (an equally weighted flows index performs worse).

The results of estimating this model using international equity index returns are presented in Panel A of Table 2.3. This panel shows that the coefficient on our hedge fund illiquidity index is positive and significant at the 5% level for all 21 countries, generating an adjusted R^2 of between 1.3% (Finland) and 8.2% (Australia), with an average of 3.4%. Of the three competing predictor variables, the innovations to the VIX index are the next most important, with 8 out of 21 index returns having a coefficient on VIX innovations that is significant at the 5% level, and another 7 significant at the 10% level. The average adjusted R^2 from models using the VIX is 3.2%. Models based on the dividend yield, lagged returns, and hedge fund flows, perform poorly – this suggests that the total endowment of hedge fund illiquid assets, and not just the per-period capital flows, are what matter for identifying the predictive relationship.

Next, we estimate a multivariate regression model where we include all of the competitor predictor variables together with the hedge fund illiquidity index in the same in-sample predictive regression:

$$r_{i,t+1} = \alpha_i + \beta_i X_t + \gamma_i \rho_t + \epsilon_{i,t+1}, \text{ where } \epsilon_{i,t+1} \sim N(0, \sigma^2_{i,t+1}). \tag{2.2}$$

where X_t is a vector that contains all of the competitor variables. The error term, $\epsilon_{i,t+1}$, is again assumed to be normally distributed with heteroskedasticity and serial correlation. The hedge fund flows are not included due to poor results in the single variable predictive regression. The last two columns of Panel A present the results from this model with and without the hedge fund illiquidity index included, to reveal the additional explanatory power of including our measure of hedge fund illiquidity. We see that the average adjusted R^2 jumps from 3.6% to 7.1% with the inclusion of our measure. Furthermore, our measure remains positive and significant at the 5% level for all 21 country indices even after the inclusion of the competitor variables. This provides strong support that there is substantial incremental predictive power

in hedge fund illiquidity for international equity returns.

The next panel of the table conducts the same in-sample predictive analysis using return data on U.S. corporate bonds. Given the large number of individual bonds (42) in our sample, we report just summaries of those results, aggregating first by ratings class (investment grade or high yield) and then by maturity.[5] The results are presented in the second panel of Table 2.3. For investment grade bonds, the hedge fund illiquidity index is the most successful single predictor across the set of five predictor variables, generating an average adjusted R^2 of 3.1% and having a significant coefficient for 14 out of the 24 individual bonds. The next most successful is the simple lagged return on U.S. corporate bonds, which generates an average adjusted R^2 of 1.8% and is significant for 12 out of 24 bonds. For high yield bonds, many of the individual variables are significant: our hedge fund illiquidity index has the second-highest number of significant coefficients (17 out of 18, compared with 18 for the lagged value-weighted equity return).

Turning to the results from the multiple predictors regression model, we see that the hedge fund illiquidity index adds substantial predictive power beyond the four competitor variables: the adjusted R^2 rises from 4.3% to 6.6% for investment grade bonds, and from 14.6% to 19.6% for high yield bonds. The number of bonds with significant coefficients on the hedge fund illiquidity index is 14 (out of 24) for investment grade bonds and 18 out of 18 for high yield bonds. Thus, the hedge fund illiquidity index is a significant predictor of U.S. corporate bond returns and appears particularly useful for high yield bonds.

The lower part of Panel B of Table 2.3 presents the results across bond maturity. We find that the results are robust across both the short and long ends of the yield curve, with a slight increase in predictive ability as measured by the adjusted R^2 for the longest maturity instruments. We explore the cross-sectional variation in predictive ability along both the maturity and the ratings dimensions further as

[5]Results for the individual bonds can be found in Appendix 2.D.

Table 2.3: In-Sample Predictive Performance

For international equities and currencies, reported are the signs of the predictors (in the case of multiple predictors, it stands for the sign of the hedge fund illiquidity index), the adjusted R^2 of the in-sample predictive regression in %, and the significance of the hedge fund illiquidity index (* for 10% and ** for 5%). The last row shows the average adjusted R^2, the total number of indices for which the coefficient of the predictor is significant (at the 10% level) and positive, and the total number of indices for which the coefficient of the predictor is significant and negative (in the case of multiple predictors, only the coefficients of the hedge fund illiquidity index are considered). For U.S. corporate bonds, the number of portfolios for which the hedge fund illiquidity index is significant at the 10% level and positive or negative is reported, and the adjusted R^2 is averaged across all portfolios. For all three asset classes, the AUM weighted flows are not included in the multiple predictor regression. For international equities and currencies, the time series starts in January 1995 and ends in December 2011. For U.S. corporate bonds, the time series starts in January 1997 and ends in December 2011. Newey-West standard errors are used.

Panel A: International Equities

	Single Predictor					**Multiple Predictors**	
Country	**Hedge Fund Illiq. Index**	**Dividend Yield**	**Lagged Returns**	**VIX AR(2) Shocks**	**Hedge Fund Flows**	**Excl. HF Illiq. Index**	**Incl. HF Illiq. Index**
Australia	(+) 8.220**	(+) 2.302	(+) 0.642	(-) 1.289	(+) -0.437	3.158	(+) 10.746**
Austria	(+) 3.954**	(-) -0.469	(+) 5.932*	(-) 8.126**	(+) 0.429	8.252	(+) 12.239**
Belgium	(+) 2.608**	(-) -0.113	(+) 7.404**	(-) 7.597*	(+) -0.001	10.228	(+) 12.617**
Canada	(+) 3.057**	(+) -0.250	(+) 1.437	(-) 4.875*	(-) -0.484	4.058	(+) 7.526**
Denmark	(+) 2.920**	(+) 2.982**	(+) 1.142	(-) 6.277**	(+) 0.019	7.945	(+) 11.129**
Finland	(+) 1.327**	(+) 0.007	(+) 3.342**	(-) -0.482	(-) -0.487	3.600	(+) 4.414**
France	(+) 2.148**	(-) -0.298	(+) 0.778	(-) 2.175*	(+) -0.176	1.508	(+) 3.947**
Germany	(+) 1.611**	(-) -0.433	(+) -0.119	(-) 1.679*	(+) -0.001	1.078	(+) 2.990**
Hong Kong	(+) 4.638**	(+) -0.223	(+) 0.825	(-) 0.941	(-) -0.072	0.467	(+) 4.806**
Ireland	(+) 2.741**	(+) -0.331	(+) 4.893*	(-) 3.032*	(+) -0.269	4.280	(+) 6.583**
Italy	(+) 1.632**	(-) -0.015	(+) -0.238	(-) 1.967*	(+) -0.232	1.789	(+) 3.897**
Japan	(+) 2.745**	(+) -0.414	(+) 2.673**	(-) 3.082**	(-) -0.374	3.483	(+) 5.881**
Netherlands	(+) 3.255**	(+) -0.480	(+) 1.716	(-) 5.669**	(+) -0.450	4.876	(+) 8.828**
New Zealand	(+) 4.722**	(+) -0.449	(-) -0.461	(-) 4.000**	(+) -0.497	5.158	(+) 11.606**
Norway	(+) 4.509**	(-) -0.152	(+) 1.199	(-) 4.419**	(+) -0.407	4.044	(+) 9.142**
Singapore	(+) 6.560**	(+) -0.494	(+) 1.009	(-) 1.595	(-) -0.418	0.832	(+) 7.960**
Spain	(+) 3.836**	(-) -0.497	(+) 0.793	(-) 1.343*	(+) -0.492	0.463	(+) 4.443**
Sweden	(+) 4.022**	(-) -0.404	(+) 1.306	(-) 1.393	(-) -0.305	1.144	(+) 5.270**
Switzerland	(+) 1.739**	(+) -0.190	(+) 1.537*	(-) 4.250**	(+) -0.105	3.381	(+) 5.217**
U.K.	(+) 4.025**	(+) 1.528*	(+) 3.489*	(-) 3.978**	(+) -0.458	5.167	(+) 8.396**
U.S.	(+) 1.583**	(+) 0.866*	(+) 0.607	(-) 0.892	(+) -0.128	1.077	(+) 2.148**
Across All (21)	**3.422 (21/0)**	**0.118 (3/0)**	**1.900 (7/0)**	**3.243 (0/15)**	**-0.255 (0/0)**	**3.618**	**7.133 (21/0)**

Panel B: U.S. Corporate Bonds

	Single Predictor						Multiple Predictors	
Rating / Maturity (# of Portfolios)	**Hedge Fund Illiq. Index**	**Lagged Returns**	**Pastor-Stamb. Trad. Liq. Factor**	**VIX AR(2) Shocks**	**VWM Excess Return**	**Hedge Fund Flows**	**Excl. HF Illiq. Index**	**Incl. HF Illiq. Index**
Inv Grade (24)	3.099 (14/0)	1.816 (12/0)	0.893 (0/8)	1.028 (6/1)	1.352 (0/8)	0.256 (0/1)	4.276	6.591 (14/0)
High Yield (18)	6.364 (17/0)	8.023 (14/0)	-0.013 (0/0)	9.741 (0/17)	7.268 (18/0)	0.487 (0/3)	14.596	19.646 (18/0)
1-3Y (7)	3.886 (5/0)	4.205 (5/0)	0.241 (0/0)	4.281 (1/4)	2.666 (3/1)	1.140 (0/2)	8.786	11.862 (5/0)
3-5Y (7)	4.995 (5/0)	6.446 (6/0)	0.294 (0/1)	5.715 (1/3)	4.893 (3/1)	0.298 (0/0)	10.302	14.002 (5/0)
5-7Y (7)	4.865 (5/0)	5.658 (5/0)	0.167 (0/1)	6.171 (1/3)	4.925 (3/1)	0.048 (0/0)	10.380	14.021 (5/0)
7-10Y (7)	4.911 (6/0)	3.978 (5/0)	0.709 (0/2)	4.168 (1/2)	3.750 (3/2)	0.206 (0/1)	8.137	11.785 (6/0)
10-15Y (7)	3.244 (4/0)	3.554 (3/0)	0.177 (0/0)	4.140 (1/3)	3.647 (3/2)	0.326 (0/1)	7.205	10.001 (5/0)
15+Y (7)	5.088 (6/0)	3.017 (2/0)	0.144 (0/4)	4.097 (1/3)	3.444 (3/1)	0.111 (0/0)	7.381	11.445 (6/0)
Across All (42)	**4.498 (31/0)**	**4.476 (26/0)**	**0.505 (0/8)**	**4.762 (6/18)**	**3.887 (18/8)**	**0.355 (0/4)**	**8.698**	**12.186 (32/0)**

Panel C: Currencies							
	Single Predictor					Multiple Predictors	
Currencies	**Hedge Fund Illiq. Index**	**Inflation Differential**	**Interest Rate Differential**	**Lagged Returns**	**Hedge Fund Flows**	**Excl. HF Illiq. Index**	**Incl. HF Illiq. Index**
Australia	(+) 5.885**	(+) 3.316**	(+) -0.412	(+) -0.042	(-) -0.362	2.923	(+) 7.187**
Canada	(+) 4.111**	(+) 3.681**	(-) -0.489	(-) -0.120	(-) -0.330	3.362	(+) 5.230**
Euro	(+) 1.406**	(+) 6.148**	(+) 0.303	(+) -0.415	(+) -0.434	5.711	(+) 5.529
Japan	(+) -0.367	(+) -0.119	(+) 1.448**	(+) -0.493	(-) -0.062	1.126	(-) 0.633
New Zealand	(+) 4.565**	(+) 4.872**	(-) 0.065	(+) -0.180	(+) -0.414	5.473	(+) 7.086**
Norway	(+) 1.883**	(+) 0.065	(+) -0.023	(+) -0.333	(+) -0.344	-0.166	(+) 1.449**
Sweden	(+) 1.468**	(+) 3.832**	(+) -0.233	(+) -0.230	(+) -0.398	4.047	(+) 4.698*
Switzerland	(+) -0.207	(+) 4.188**	(+) 2.093**	(-) -0.462	(+) -0.473	5.047	(-) 4.713
U.K.	(+) -0.001	(+) 0.132	(-) 0.964	(+) -0.270	(+) -0.142	0.676	(+) 0.465
Across All (9)	**2.083 (6/0)**	**2.902 (6/0)**	**0.413 (2/0)**	**-0.283 (0/0)**	**-0.329 (0/0)**	**3.133**	**4.110 (5/0)**

part of our empirical tests of the predictions of the equilibrium model presented in the next section.

Panel C of Table 2.3 presents the results of in-sample regressions for currency returns. As is well known (see, for example, Meese and Rogoff (1983)), currencies are generally very hard to predict, and this is confirmed by our analysis. However, our hedge fund illiquidity index performs well: in univariate regressions, it generates an average adjusted R^2 of 2.1% (slightly lower than that for the inflation differential, with 2.9%) and is statistically significant for six out of the nine currencies under consideration. In the multiple predictors regression results, we see that the adjusted R^2 increases from 3.1% to 4.1% with the inclusion of the hedge fund illiquidity index, and the variable is significant for five out of the nine currencies. Thus even for the notoriously challenging task of predicting currency returns, our hedge fund illiquidity index adds value.

Figures 2.2 and 2.3 show the length of time over which the hedge fund illiquidity index retains predictive ability for future asset returns. Figure 2.2 shows the results for the single variable predictive regressions and Figure 2.3 for the multiple variables predictive regressions. The line in each figure depicts the average adjusted R^2, while the bars show the number of assets for which the hedge fund illiquidity index is statistically significant. Along the horizontal axis, we vary the forecast horizon: $h = 1$ corresponds to our baseline regression where we forecast returns using the one-month lagged illiquidity index, and $h = 6$ corresponds to forecasting returns using six-month lagged hedge fund illiquidity. The figure shows that the forecasting power of the hedge fund illiquidity index is quite long-lasting: the forecasting performance for all three asset classes only deteriorates significantly at the seven to eight month horizons.[6] Given our motivation for considering this variable, namely, the illiquidity

[6]We note here that the hedge fund illiquidity index is persistent, with first-order autocorrelation of 0.84, and a 95% confidence interval of [0.72,0.96]. However, an augmented Dickey-Fuller test for a unit root very strongly rejects the null in favor of stationarity, and the point estimate suggests that the persistence should not cause us any real problems with inference in forecasting. Thus, this persistence in predictive power is not attributable to excessive persistence in our illiquidity index.

of hedge fund portfolios as a signal of their ability to provide liquidity to markets, this long-lasting forecasting power is not particularly surprising – it could easily take up to six months for an illiquid hedge fund to unwind its portfolio and return to the role of liquidity provision to asset markets. We motivate this insight more rigorously in our discussion of our model in Section 2.4.

2.3.2 The out-of-sample predictive power of hedge fund illiquidity

The above analyses used the full sample of data to estimate the relationship between hedge fund illiquidity and asset returns. In this section, we consider the out-of-sample performance of models based on hedge fund illiquidity as well as competitor predictor variables. Out-of-sample evaluation is increasingly an important metric for predictive models of asset returns, on account of its greater relevance for real-world investors (see, for example, Campbell and Thompson (2008) and Goyal and Welch (2008), for analyses of the predictive ability of various models for the equity premium). In this sub-section, we take the perspective of a real-world investor engaging in forecasting the set of asset returns introduced earlier, using our index of hedge fund illiquidity.

Given the well-known difficulty of out-of-sample asset return prediction with multiple variables, we focus on simple single-variable predictive models, as in equation (2.1). We estimate the parameters of the models using a rolling window of 60 months of data, so our in-sample period is January 1995 to December 1999, and our out-of-sample period is January 2000 to December 2011, for international equities and currencies. For U.S. corporate bonds, the out-of-sample period is January 2002 to December 2011. We report the OOS R^2 from each of the models, and we extend the Clark and West (2006) test to enable formal comparison of the above model with a model that only includes a constant. The original Clark-West test requires that the smaller model contains no parameters at all, whereas for asset returns, it

Figure 2.2: In-Sample Single Variable Predictability at Different Horizons

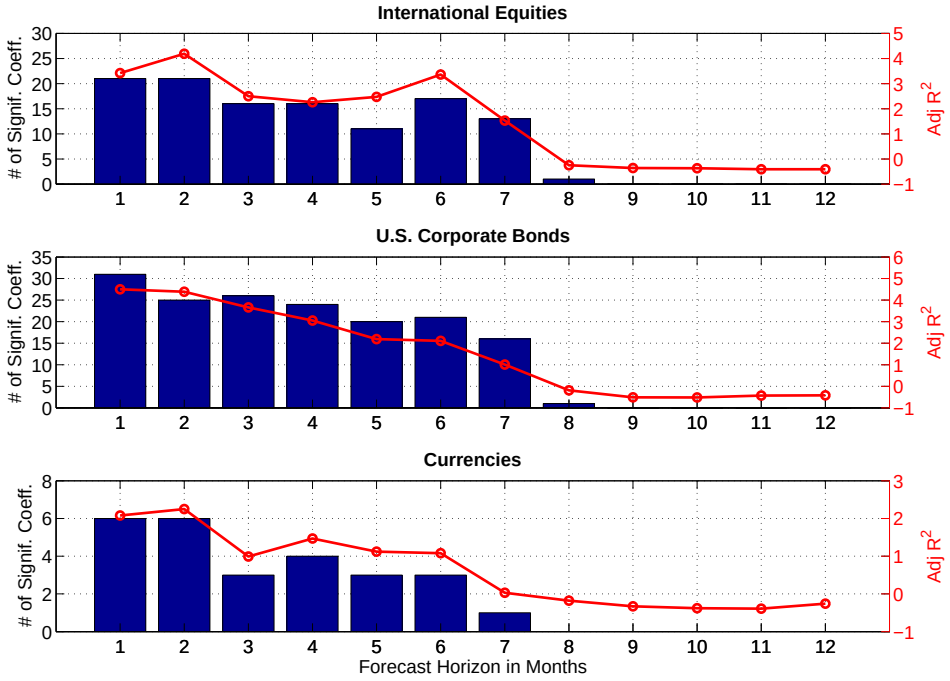
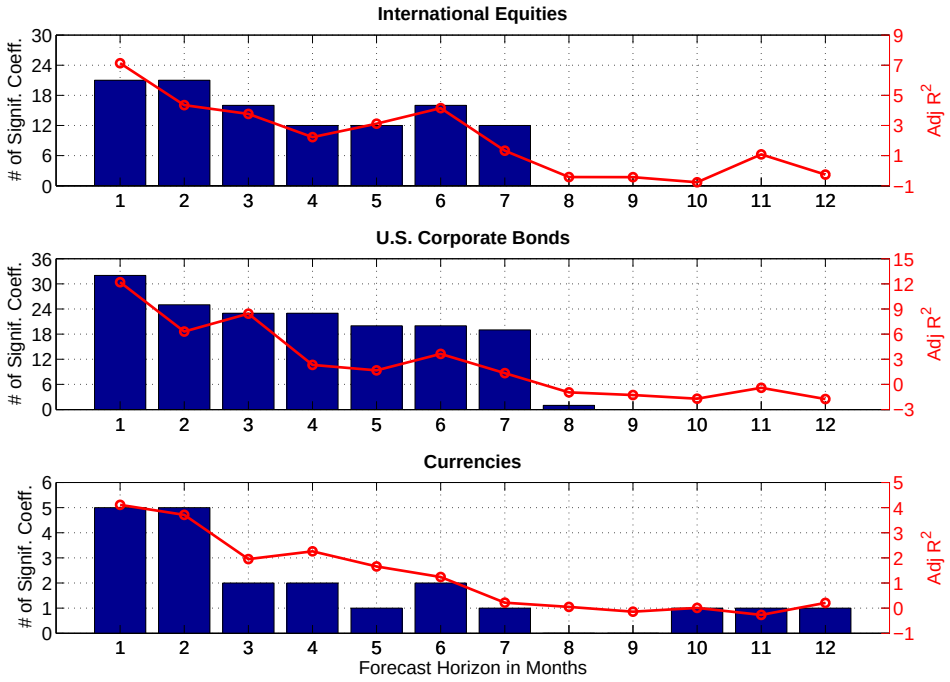


Figure 2.3: In-Sample Multiple Variables Predictability at Different Horizons



is more reasonable to take a model with just a constant term as a benchmark. The details of this simple extension are presented in the Appendix 2.A. There, we also show its finite sample performance via a small simulation study.

Panel A of Table 2.4 presents the out-of-sample results for predicting international equity index returns. Forecasts based on our hedge fund illiquidity index generate an average OOS R^2 of 3.2%, ranging from -0.9% for Japan to 9.0% for Australia. In only a single case (Japan) is the OOS R^2 negative, and for 20 out of the 21 countries the extension of the Clark-West test rejects the constant model in favour of one which includes the hedge fund illiquidity index.[7] In contrast, the out-of-sample forecasts based on the three competing variables (dividend yield, lagged returns, and shocks to the VIX index) generate negative average OOS R^2, and are significant for a maximum of four countries in the sample. These results provide support for the result obtained in our full-sample analysis: hedge fund illiquidity is useful for predicting international equity index returns.

In Panel B of Table 2.4, we find similarly good results for U.S. corporate bonds. The average OOS R^2 from forecasts based on our hedge fund illiquidity index is 4.4%, and is positive for all ratings and maturity subsets of these bonds. Further, we formally reject the constant model in favor of our model for 28 out of the 42 bonds. The forecasts based on the competing variables (lagged returns, shocks to the VIX index, and the lagged market return) perform much worse: all three generate negative average OOS R^2. The forecasts based on the lagged market return perform the best of these three variables, with an average OOS R^2 of -2.8%, and significantly beating the benchmark model for 17 out of 42 bonds. For the lagged market return, the only subset of bonds for which the average OOS R^2 is positive is the high yield group, and even there the performance of this predictor variable is inferior to that

[7]Note that a higher OOS R^2 does not necessarily translate to greater significance in the Clark-West test. For example, the forecasts based on the hedge fund illiquidity index have an OOS R^2 of 3.7% for the Netherlands and reject the null that the constant model is best at the 10% level, while for Italy, the OOS R^2 is 2.3%, and we reject the null in favor of the larger model at the 5% level.

obtained when using our index of hedge fund illiquidity.

Panel C of Table 2.4 presents the results for out-of-sample currency return prediction. As might be expected given the well-documented difficulty of predicting currency returns, the OOS R^2 is lower for this asset class, and indeed for all four variables, we find a negative average OOS R^2, with the highest (-0.03%) coming from our hedge fund illiquidity index. For three of these currencies, we are able to reject the constant model. These currencies are the Australian Dollar, the Canadian Dollar, and the New Zealand Dollar. The forecasts for the Norwegian Krona have a positive OOS R^2, but we do not reject the null. These are all relatively high interest-rate commodity currencies – which are generally associated with the long side of the frequently studied carry trade (see Burnside, Eichenbaum, and Rebelo (2011), for an overview of the carry trade literature). This finding is explored further when we test our model predictions below. Finally, when comparing the performance of the illiquidity index to forecasts of currency returns based on competing variables, the primary significant result is the usefulness of the inflation differential for predicting the euro return: this forecast generates an OOS R^2 of 4.8% and is significantly better than the constant model. The interest rate differential rejects the constant model for New Zealand and Sweden. The lagged returns fail to generate a positive OOS R^2 or reject the historical mean return model for any of the currencies.

It is worth noting here that our out-of-sample period, 2000 to 2011, spans the financial crisis and the Great Recession. Unfortunately, hedge fund data only became available in the early to mid 1990s, and so we are unable to conduct formal analysis of the impact of recessions on the predictive relationships we document here. To gain some insight into the impact, we constructed a time series of the cumulative sum of squared errors of the historical mean return model and the predictive regression, similar to that used in Goyal and Welch (2008). For all three asset classes, we find that our hedge fund illiquidity index performs particularly well during the financial crisis.

Table 2.4: Out-of-Sample Forecast Performance

Reported are the results of rolling 5 year OOS forecasts with single predictor regressions. The dependent variable is the monthly log excess return for international equities and U.S. corporate bonds and the monthly log return for currencies. For international equities and currencies, we report the OOS R^2 and whether we can reject the historical average model based on an extension of the Clark-West test-statistic, which is denoted by * for 10% significance and ** for 5% significance. For U.S. corporate bonds, the number of portfolios for which we can reject the historical average model at the 10% level is in parentheses. The last row in every panel shows the average OOS R^2 and the total number of portfolios for which the historical average model is rejected. For international equities and currencies, the time series starts in January 1995 and ends in December 2011. The time series for U.S. corporate bonds starts in January 1997 and ends in December 2011.

Panel A: International Equities				
Country	**Hedge Fund Illiq. Index**	**Dividend Yield**	**Lagged Returns**	**VIX AR(2) Shocks**
Australia	9.039**	-3.051	-1.423	-7.714
Austria	4.366**	-0.986*	1.006	-1.871*
Belgium	3.260**	-2.540	5.024*	-1.187*
Canada	2.711**	-4.085	-3.368	-2.121
Denmark	2.680**	0.064	0.240	-3.062*
Finland	1.815**	-4.952	1.623**	-3.190
France	2.296**	-8.358	-1.619	-4.776
Germany	0.720*	-17.077*	-2.637	-3.694
Hong Kong	3.328**	-5.041	-1.155	-8.238
Ireland	2.739*	-10.150	7.579*	0.027
Italy	2.311**	-3.413	-1.336	-3.742
Japan	-0.887*	-2.384	0.533*	-0.360**
Netherlands	3.679*	-8.505	-2.661	-5.643
New Zealand	5.764**	-46.369	-2.256	-0.818
Norway	4.836**	-4.737	-2.499	-1.102
Singapore	3.973**	-7.015	-2.210	-7.410
Spain	4.569**	-4.332	-2.177	-5.448
Sweden	3.119**	-11.399	-0.857	-7.212
Switzerland	0.963	-4.609	0.200	-0.155
U.K.	4.653**	-6.286	1.520	-2.912
U.S.	1.657*	-1.283	-1.300	-5.120
Across All (21)	**3.219 (20)**	**-7.453 (2)**	**-0.370 (4)**	**-3.596 (4)**

Panel B: U.S. Corporate Bonds

Rating/Maturity	Hedge Fund Illiq. Index	Lagged Returns	VIX AR(2) Shocks	VWM Ex. Return
Inv Grade (24 Portfolios)	3.300 (11)	-7.016 (0)	-8.238 (2)	-7.392 (0)
High Yield (18 Portfolios)	5.835 (17)	-4.076 (10)	0.236 (3)	3.259 (17)
1-3Y (7 Portfolios)	4.343 (5)	-17.113 (1)	-3.332 (1)	-3.591 (2)
3-5Y (7 Portfolios)	5.053 (4)	-2.460 (3)	-3.417 (1)	-1.658 (3)
5-7Y (7 Portfolios)	4.664 (5)	-3.829 (1)	-5.024 (2)	-2.417 (3)
7-10Y (7 Portfolios)	4.594 (5)	-5.427 (1)	-7.021 (0)	-3.671 (3)
10-15Y (7 Portfolios)	3.059 (4)	-0.438 (1)	-5.100 (0)	-2.803 (3)
15+Y (7 Portfolios)	4.604 (5)	-5.268 (3)	-3.743 (1)	-2.824 (3)
Across All (42)	**4.386 (28)**	**-5.756 (10)**	**-4.606 (5)**	**-2.827 (17)**

Panel C: Currencies

Currency	Hedge Fund Illiq. Index	Inflation Diff.	Interest Rate Diff.	Lagged Returns
Australia	4.098**	-0.342	-0.631	-2.847
Canada	1.372**	0.744*	0.875	-2.364
Euro	-0.405	4.786**	-1.682	-2.663
Japan	-4.229	-2.231	-6.230	-4.658
New Zealand	3.545**	1.143*	1.770*	-1.910
Norway	0.790	-2.565	-1.268	-4.893
Sweden	-1.365	1.012	1.067**	-3.304
Switzerland	-2.186	-0.109*	-2.174	-2.316
U.K.	-1.895	-2.905	0.196	-1.250
Across All (9)	**-0.030 (3)**	**-0.052 (4)**	**-0.897 (2)**	**-2.912 (0)**

We view this as consistent with our explanation for the source of the predictive power: during this period, capital constraints on hedge fund portfolios ("funding liquidity" a la Gromb and Vayanos (2002) and Brunnermeier and Pedersen (2008)) were probably most binding, translating into constraints on hedge funds' liquidity provision to asset markets.

The next section presents a simple model which provides a more rigorous foundation for our results and generates additional empirical predictions which we subsequently test.

2.4 A Model of Hedge Fund Illiquidity and Asset Return Determination

Our model is based on the limits of arbitrage framework laid out in Gromb and Vayanos (2010). In our simple model, there is one risky asset, which is traded by a risk-averse hedge fund and by noise traders. The hedge fund effectively acts as the market maker for the risky asset, which is subject to demand shocks from noise traders.

2.4.1 The timeline of the model

The model contains three periods. In period 0, the hedge fund in our model inherits an illiquid endowment θ, which we assume it cannot sell, and a cash endowment C_0. For simplicity, the return on both the illiquid endowment and the cash endowment are set to zero. In period 1, the hedge fund determines its demand for the risky asset as a function of price. The net demand of the noise traders is then realised and trading occurs, which determines the equilibrium price and quantity traded for the risky asset. In period 2, the dividend on the risky asset is realised.

In the next section we describe the hedge fund's optimization problem, and then we find an expression for the equilibrium price and quantity of the risky asset traded.

2.4.2 The hedge fund's objective function

The objective function of the hedge fund is assumed to take a simple quadratic form, and the fund is subject to one additional, important constraint. In our model, the hedge fund is concerned about the illiquidity of its portfolio, because it potentially faces fund outflows due to redemptions by investors. As a result, it requires a sufficient quantity of cash to pay out investors who withdraw their funds. Investor redemptions' impact on the liquidity provision of hedge funds is also highlighted in Mitchell, Pedersen, and Pulvino (2007).

We model the illiquidity constraint in a straightforward fashion. The hedge fund has an expectation of the maximum level of fund outflows in a given time period, Φ_{max}. We set $\Phi_{max} = \phi_{max}(C_0 + \theta)$, with $\phi_{max} \in [0, 1]$. Hence, Φ_{max} is proportional to the total AUM, $C_0 + \theta$, of the hedge fund. For example, when $\phi_{\max} = 0.5$, the hedge fund expects that, in the worst case scenario, its investors will withdraw 50% of the AUM. We model a convex "shortfall penalty" if the fund does not hold sufficient liquid assets to cover $\Phi_{\max}$ after the trade in the risky asset. Letting x_1 denote the demand of the hedge fund for the risky asset in period 1, and p_1 the equilbrium price of the risky asset in period 1, the convex shortfall penalty is given by

$$\begin{array}{ccc} 0 & if & \Phi_{\max} \leq C_1 \\ \frac{1}{2}(\Phi_{\max} - C_1)^2 & if & \Phi_{\max} > C_1, \end{array} \tag{2.3}$$

where $C_1 = C_0 - x_1 p_1$ is the new level of cash that the fund holds after its purchases or sales of the risky asset.

The fund's objective function can therefore be written as

$$Q(x_1) = x_1(E[d_2] - p_1) - \frac{\alpha}{2} x_1^2 \sigma_d^2 - \lambda \frac{1}{2}(\Phi_{\max} - C_0 + x_1 p_1)^2 \mathbf{1}\{\Phi_{\max} > C_0 - x_1 p_1\}. \tag{2.4}$$

Here, d_2 is the dividend, which has variance σ_d^2, α is the risk aversion of the fund, and the strictly positive λ is the weight on the illiquidity constraint in the fund's

objective function.

The first order condition of the hedge fund's maximization problem for $\Phi_{\max} \leq C_1$ is given by

$$E[d_2] - p_1 - \alpha x_1 \sigma_d^2 = 0. \tag{2.5}$$

This results in the mean-variance demand function:

$$x_1 = \frac{E[d_2] - p_1}{\alpha \sigma_d^2}. \tag{2.6}$$

In the case where $\Phi_{\max} > C_1$, the first order condition is

$$E[d_2] - p_1 - \alpha x_1 \sigma_d^2 - \lambda(\Phi_{\max} - C_0 + x_1 p_1) p_1 = 0, \tag{2.7}$$

and the demand function becomes

$$x_1 = \frac{E[d_2] - p_1 - \lambda(\Phi_{\max} - C_0) p_1}{\alpha \sigma_d^2 + \lambda p_1^2}. \tag{2.8}$$

The second demand function reveals that the lower the value of C_0 and the higher the illiquid endowment θ (which raises $\Phi_{\max}$), the lower the quantity of the risky asset x_1 which the hedge fund will hold. This yields an important prediction: an illiquid hedge fund is more willing to sell the risky asset than to buy it. This feature of the model, which we explain more fully below, is what underpins the model's explanation for the predictive power of our hedge fund illiquidity index for asset returns.

2.4.3 Market clearing and the equilibrium price

The demand of the noise traders, u_1, is an exogenous shock, and determines the price that the hedge fund demands to absorb this shock in equilibrium. To compute the equilibrium price, we need to clear markets. As in Gromb and Vayanos (2010),

we interpret the demand shock u_1 as net aggregate demand, which implies that the risky asset is in zero net supply. Hence, the market clearing condition for a hedge fund which holds enough cash to cover potential fund outflows takes the form

$$\frac{E[d_2] - p_1}{\alpha\sigma_d^2} + u_1 = 0. \tag{2.9}$$

The equilibrium price is then given by

$$p_1 = E[d_2] + u_1\alpha\sigma_d^2. \tag{2.10}$$

For the case where $\Phi_{\max} > C_1$, the market clearing condition is

$$\frac{E[d_2] - p_1 - \lambda(\Phi_{\max} - C_0)p_1}{\alpha\sigma_d^2 + \lambda p_1^2} + u_1 = 0, \tag{2.11}$$

which results in the equilibrium price:[8]

$$p_1 = \frac{(1 + \lambda(\Phi_{\max} - C_0)) - \sqrt{(1 + \lambda(\Phi_{\max} - C_0))^2 - 4(\lambda u_1)(E[d_2] + \alpha\sigma_d^2 u_1)}}{2\lambda u_1}. \tag{2.12}$$

2.4.4 Model-implied hedge fund illiquidity and asset returns

We measure the model-implied illiquidity level of the hedge fund's portfolio before the trade with the noise traders as the ratio of the value of its illiquid endowment to the value of its total assets under management:

$$\rho^{\text{model}} = \frac{\theta}{(\theta + C_0)}. \tag{2.13}$$

[8]The quadratic equation for the equilibrium price has a second root. However, we ignore the second root, as the implied equilibrium price is not sensible regardless of the parameter values. In our simulation, when the hedge fund buys, the second root price is negative and generally smaller than -100. When the hedge fund sells, the second root price is positive and generally greater than 100.

Hence, the lower C_0 relative to θ, the greater the illiquidity of the hedge fund's portfolio.

We next specify parameters and simulate the model in order better to understand its predictions. $\phi_{\max}$ is determined endogenously. We set $\phi_{\max} = (\theta/(\theta + C_0))^2 = \rho^2$. The intuition is that the higher the ratio of the illiquid endowment to the total AUM, the more concerned are investors about the illiquidity of the hedge fund. Due to stronger illiquidity concerns, redemptions are expected to be higher. We assume a convex relationship between portfolio illiquidity and redemptions, although a linear relationship would lead to similar results. Panel A of Table 2.5 shows the specific parameter values that we use. The two sources of uncertainty in our model, d_2 and u_1, are drawn from Normal distributions with means $\mathrm{E}[d_2] = 1$ and $\mathrm{E}[u_1] = 0$, and variances $\sigma_d^2 = 0.02$ and $\sigma_u^2 = 0.25$, respectively. λ, the weight on the shortfall penalty in the hedge fund's objective function, is set to 0.005, α, the risk-aversion of the hedge fund, to 3, and θ, the initial endowment of the illiquid asset, to 10.

Whether the shortfall penalty affects the utility of the hedge fund depends on p_1, but p_1 depends on the shortfall penalty. Because of this codependency, we need a numerical procedure to simulate the model. The numerical procedure is given in Appendix 2.C.

Panel B of Table 2.5 looks at comparative statics as we vary C_0, which determines the level of illiquidity of the fund. In this exercise, values for d_2 and u_1 are drawn from their distributions whose parameters are given in Panel A, while the value of C_0 is varied across different simulations. We simulate the two period model $10,000$ times for each value of C_0 and compute the average expected return, ρ^{model}, and the "price discounts" (relative to the price when $\lambda = 0$) at which noise traders buy from and sell to the hedge fund.

The top row in this panel simply sets the liquidity constraint on the fund to zero, $\lambda = 0$, as a benchmark, so the values of C_0 and $\phi_{\max}$ are not relevant. The expected return is computed as the average realisation of $\frac{1}{p_1}$. The noise traders buy

Table 2.5: Model Simulation

Panel A reports the parameter values used for the model simulations in Panel B and C. Panel B reports the statistics for model simulations based on fixed values of C_0. For each value of C_0, the model is simulated $10,000$ times. The parameters d_2 and u_1 are drawn randomly from Normal distributions with means of 1 and 0 (variances 0.02 and 0.25), respectively, for each of the $10,000$ simulations. The ρ, returns, and prices, are averaged across the $10,000$ simulations. Panel C reports the estimates of α and γ for the predictive regression $r_{t+1} = \alpha + \gamma\rho_t + \epsilon_{t+1}$ estimated with $10,000$ data points of simulated data from our model. The values for d_2, u_1, and C_0, are drawn randomly for each of the $10,000$ simulations. The values for d_2 and u_1 are drawn from the same distributions as in Panel B. C_0 is drawn from a uniform distribution with a lower bound of 0 and an upper bound of 8.

Panel A: Parameter Values

$E[d_2]$	σ_d^2	λ	α	θ	$E[u_1]$	σ_u^2
1	0.02	0.005	3	10	0	0.25

Panel B: Model Simulation for Fixed C_0

No Liquidity Constraint						
C_0	C_0/θ	$\phi_{max}(\theta + C_0)$	ρ^{model}	**Expected Return**	**Noise Traders Buy Price**	**Noise Traders Sell Price**
NA	NA	NA	NA	0.085%	1.024	0.976

With Liquidity Constraint						
				Expected	**Price Discount**	
C_0	C_0/θ	$\phi_{max}(\theta + C_0)$	ρ^{model}	**Return**	**NT Buy**	**NT Sell**
5.0	0.5	6.7	0.667	0.918%	0.619%	-1.009%
4.0	0.4	7.1	0.714	1.656%	1.344%	-1.726%
3.0	0.3	7.7	0.769	2.432%	2.094%	-2.467%

Panel C: Predictive Regression on Simulated Data

	Estimates		
	All Observations	**Noise Traders Buy**	**Noise Traders Sell**
α	-0.071	-0.071	-0.073
γ	0.125	0.091	0.160

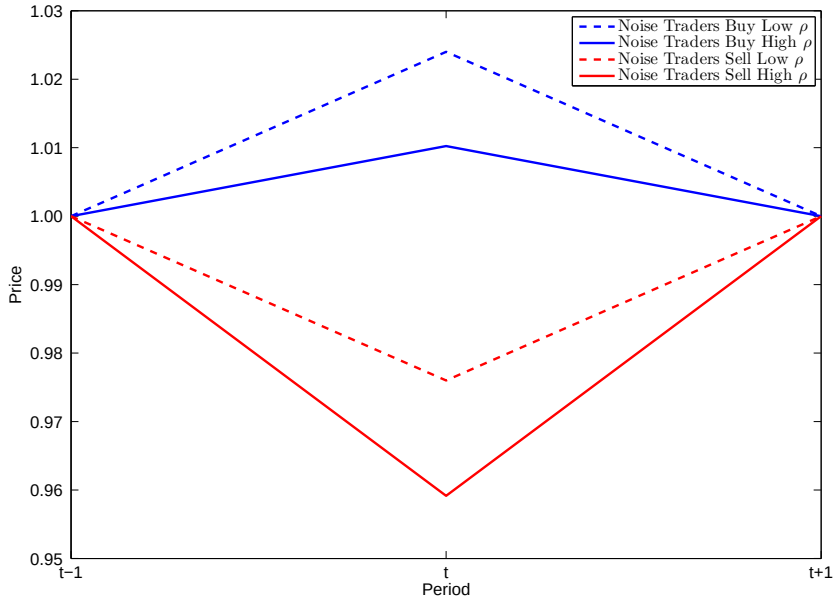
from the hedge fund at a price slightly above the expected fundamental value and sell to the hedge fund at a slightly lower price than the expected fundamental value as compensation for the hedge fund's market-making services.

As we move down the rows of Panel B, we impose the liquidity constraint, and vary the level of starting cash, C_0, resulting in a cash-to-liquid assets ratio of 50%, 40%, and 30%, respectively, for each of the bottom three rows. $\Phi_{\max} = \phi_{\max}(\theta + C_0)$ varies as we vary the level of cash. The ρ^{model} values that these correspond to are shown in the next column, and we can see that a decrease in C_0 leads to a higher average ρ^{model} as well as, in the next column, to a higher average expected return. This expected return is higher than the return in the benchmark case, which reflects the extra compensation needed to convince the illiquid hedge fund to purchase more of the risky asset, as well as the price discount offered by the hedge fund to the noise traders when they purchase the risky asset. While the illiquid hedge fund is willing to sell the risky asset for a lower price in the face of a positive noise trader demand shock, in order to increase the liquidity of its portfolio and avoid hitting the shortfall penalty, the fund is reluctant to buy the risky asset and requires more compensation to do so. Relative to the case where $\lambda = 0$, the price discounts for the noise traders are substantial – reaching roughly 2.0% when the hedge fund discounts the price to sell to noise traders, and roughly -2.5% when the hedge fund demands compensation in order to accept more of the risky asset into its portfolio, when $C_0 = 3$.

Finally, we again simulate the model $10,000$ times but with C_0 drawn randomly for each simulation. C_0 is drawn from a uniform distribution with a lower bound of 0 and an upper bound of 8, and d_2 and u_1 drawn as described above. We take the $10,000$ observations of ρ^{model} and the realised returns in the second period, r_2, derived from these simulations and regress:

$$r_2^{\text{realised}} = \alpha + \gamma \rho^{\text{model}} + \epsilon_2, \text{ where } \epsilon_2 \sim N(0, \sigma^2). \tag{2.14}$$

Figure 2.4: Model Implied Return Reversal



This figure depicts the average return reversal implied by the model for the parameter values given in Table 2.5. The high ρ case corresponds to $C_0 = 4$, and the low ρ case corresponds to $\lambda = 0$.

The model-implied α and γ are reported in the first column of Panel C of Table 2.5. We also separately estimate α and γ conditional on the sign of realised u_1, corresponding to realisations in which noise traders buy and noise traders sell. This reveals that the coefficient γ is positive regardless of the sign of the shock, which is consistent with the intuition of the model, and a feature that we confirm in our auxiliary empirical tests below.

Figure 2.4 provides a graphical representation of the model predictions – the figure shows the comparative static of the price path compared across levels of illiquidity ρ^{model} of the hedge fund. Compared to a liquid hedge fund, an illiquid hedge fund buys for a *lower* price when noise traders sell, which implies a *greater* return reversal. It also sells for a *lower* price when noise traders buy the risky asset, implying a *smaller* return reversal (in absolute values). When these predictions are taken together, the model implies that a high ρ^{model} predicts high returns.

2.4.5 Empirical predictions of the model

This section presents tests of predictions from the model described above, using empirical data. Above, we show that, regardless of the sign of the noise trader shock, the model predicts that the coefficient on ρ will be positive. To test this, we use an approach from Pastor and Stambaugh (2003), who identify positive and negative noise trader demand shocks using the sign of the lagged return. We thus estimate a specification which includes the competitors for each asset class (including the lagged returns), as well as a dummy for the sign of the lagged return that we interact with our hedge fund illiquidity index:

$$r_{i,t+1} = \alpha_i + \gamma^{\text{NT Sell}} \rho_t I_{\{r_{i,t}<0\}} + \gamma^{\text{NT Buy}} \rho_t I_{\{r_{i,t}>0\}} + \beta X_{i,t} + \epsilon_{i,t+1}, \tag{2.15}$$

where $\epsilon_{i,t+1} \sim N(0, \sigma_i^2)$. The prediction of our theoretical model is that both $\gamma^{\text{NT Sell}}$ and $\gamma^{\text{NT Buy}}$ will be positive. To increase the precision of the parameter estimates, we estimate the specification above as a panel, grouping all assets i within each of the three broad asset classes. We allow for asset-specific fixed effects in the estimation, and we impose that the slope coefficients are the same across all assets within each class. In all panel regressions, standard errors are adjusted for contemporaneous correlation using a delete-cross-section jackknife method (see Shao and Wu (1989) and Shao (1989) for examples of the implementation of the jackknife estimator).

Equation (2.15) nests our specification from Section 2.3.1, as when $\gamma^{\text{NT Sell}} = \gamma^{\text{NT Buy}}$, the indicator variable effectively drops out. We consider the full model, and also restricted versions where we include just one or the other of the interaction terms. Table 2.6 presents the results for the three asset classes. In all three asset classes, corresponding to the three panels in the table, we see that the coefficients $\gamma^{\text{NT Sell}}$ and $\gamma^{\text{NT Buy}}$ are estimated to be positive and statistically significant (with the one exception of currencies, for which $\gamma^{\text{NT Buy}}$ is positive, but not statistically significant) when including them both in the regression.

Table 2.6: Panel Estimation Conditioning on Sign of Lagged Return

Reported are the results of a fixed effects panel estimation with jackknife standard errors clustered by time. The dependent variables are monthly log excess returns or monthly log returns in the case of currencies. The independent variables are lagged by one month and divided by their respective standard deviations. For international equities and currencies, the time series starts in January 1995 and ends in December 2011. The time series for U.S. corporate bonds starts in January 1997 and ends in December 2011. Estimates significant at the 10% level are denoted by *, and estimates significant at the 5% level are denoted by **.

Panel A: International Equities	(1)	(2)	(3)
Dividend Yield	0.038	0.035	-0.060
	(0.138)	(0.126)	(-0.228)
Lagged Returns	1.203**	-0.099	0.616
	(2.104)	(-0.198)	(1.126)
VIX AR(2) Shocks	-0.919	-1.037	-1.078*
	(-1.412)	(-1.541)	(-1.655)
HF Illiq. Index×Negative Lag Ret. Dummy	1.243**		2.108**
	(2.533)		(3.302)
HF Illiq. Index×Positive Lag Ret. Dummy		0.552	1.634**
		(1.338)	(2.881)
Adj R2 (%)	**4.950**	**3.507**	**7.280**
Panel B: U.S. Corporate Bonds			
Lagged Returns	0.582**	0.170	0.378
	(2.393)	(0.788)	(1.566)
Pastor Stambaugh Traded Liq. Factor	-0.341**	-0.311**	-0.284**
	(-2.374)	(-2.204)	(-2.063)
VIX AR(2) Shocks	-0.464	-0.553	-0.541
	(-1.332)	(-1.506)	(-1.565)
VWM Excess Returns	-0.061	-0.150	-0.133
	(-0.239)	(-0.564)	(-0.540)
HF Illiq. Index×Negative Lag Ret. Dummy	0.261		0.630**
	(1.454)		(2.767)
HF Illiq. Index×Positive Lag Ret. Dummy		0.371**	0.695**
		(2.688)	(3.839)
Adj R2 (%)	**6.378**	**6.892**	**8.889**
Panel C: Currencies			
Inflation Differential	0.515**	0.511**	0.438**
	(3.858)	(3.796)	(3.295)
Interest Rate Differential	0.034	0.018	0.029
	(0.201)	(0.101)	(0.172)
Lagged Returns	0.186	-0.103	0.071
	(0.779)	(-0.477)	(0.294)
HF Illiq. Index×Negative Lag Ret. Dummy	0.239		0.482*
	(1.323)		(1.730)
HF Illiq. Index×Positive Lag Ret. Dummy		0.181	0.438
		(1.063)	(1.615)
Adj R2 (%)	**3.033**	**2.920**	**3.678**

Although not directly implied by the model, we can expect that the model prediction – for a high ρ, the return reversal is greater when noise traders sell and smaller when noise traders buy – is more pronounced if the risky asset is more illiquid. The intuition is that an illiquid hedge fund is even more reluctant (eager) to buy (sell) the risky asset when the risky asset is more illiquid. Hence, our ρ should have more predictability for illiquid assets. To test this for each asset class, we run another fixed-effect panel regression, in which we include ρ, ρ interacted with a dummy which captures assets with greater illiquidity, and the set of all competitor predictor variables X for the asset returns in a given class. As before, the model is estimated for each asset class. The specification is

$$r_{i,t+1} = \alpha_i + \gamma\rho_t + \gamma^{\text{Illiq.}}\rho_t I_i^{\text{Illiq.}} + \beta X_{i,t} + \epsilon_{i,t+1}, \text{ where } \epsilon_{i,t+1} \sim N(0, \sigma_i^2). \quad (2.16)$$

The illiquidity dummy, $I_i^{Illiq.}$, takes the value one if portfolio i belongs to a more illiquid subgroup of the asset class. For international equities, we use the market capitalisation and the turnover average over the whole sample period as illiquidity proxies (for values below the median across all equity markets in the sample, the dummy variable takes the value one). For U.S. corporate bonds, high yield and longer maturity (over 5 years) bonds are considered more illiquid. In the case of currencies, we use the Libor interest rates and the percentage bid-ask spread averaged over the whole sample period as an indicator for illiquidity (for values above the median across all currencies in the sample, the dummy variable takes the value one).

We expect the coefficient of the illiquidity dummy to be significant and positive, and Table 2.7 shows that this is the case, consistently for all asset classes. The standard errors are again adjusted for contemporaneous correlation using the delete-cross-section jackknife method.

Overall, the simple model presented in this section helps to explain our main

Table 2.7: Panel Estimation with Illiquidity Dummy Variables

Reported are the coefficient estimates and t-stats based on jackknife standard errors clustered by time. The dependent variables are monthly log excess returns, except for currencies, where we use the monthly log return. The regressors are lagged by one month and standardised. For international equities, the illiquidity dummy depends on either the total sample average market cap or average turnover standardised by market cap (the asset class median is the benchmark). For U.S. corporate bonds, we use a high yield dummy and a long maturity dummy. For currencies, the illiquidity dummy depends either on the total sample average 1M Libor rate or average percentage spread (the asset class median is the benchmark). The time series starts in January 1995 and ends in December 2011, except for U.S. corporate bonds the start is in January 1995. 10% significance is denoted by * and 5% by **.

Panel A: International Equities	(1)	(2)	(3)
Dividend Yield	-0.065	-0.061	-0.064
	(-0.246)	(-0.233)	(-0.244)
Lagged Returns	0.113	0.111	0.109
	(0.287)	(0.282)	(0.276)
VIX AR(2)	-1.126*	-1.127*	-1.128*
	(-1.676)	(-1.677)	(-1.679)
HF Illiq. Index	1.297**	1.164**	1.179**
	(3.213)	(3.058)	(2.900)
HF Illiq. Index×(Low Mkt. Cap Dummy)		0.278**	
		(2.165)	
HF Illiq. Index×(Low Turnover Dummy)			0.248**
			(2.161)
Adj R2 (in %)	**6.826**	**6.848**	**6.839**
Panel B: U.S. Corporate Bonds			
Lagged Returns	0.339**	0.341**	0.338**
	(2.118)	(2.118)	(2.110)
Pastor Stambaugh Traded Liq. Factor	-0.282**	-0.282**	-0.282**
	(-2.043)	(-2.044)	(-2.042)
VIX AR(2) Shocks	-0.549	-0.549	-0.549
	(-1.551)	(-1.553)	(-1.551)
VWM Excess Return	-0.144	-0.144	-0.144
	(-0.570)	(-0.569)	(-0.569)
HF Illiq. Index	0.505**	0.229**	0.375**
	(3.768)	(1.995)	(3.137)
HF Illiq. Index×(High Yield Dummy)		0.666**	
		(2.626)	
HF Illiq. Index×(Maturity >5Y Dummy)			0.197**
			(2.713)
Adj R2 (in %)	**8.993**	**10.335**	**9.090**
Panel C: Currencies			
Inflation Differential	0.439**	0.437**	0.442**
	(3.240)	(3.228)	(3.258)
Interest Rate Differential	0.025	0.031	0.026
	(0.149)	(0.187)	(0.156)
HF Illiq. Index	0.323*	0.206	0.203
	(1.747)	(1.134)	(1.255)
HF Illiq. Index×(High Int. Rate Dummy)		0.264**	
		(2.463)	
HF Illiq. Index×(High Avg. Spread in %)			0.269**
			(2.229)
Adj R2 (in %)	**3.743**	**3.860**	**3.867**

predictive results in Section 2.3 and generates additional testable predictions that are confirmed in our data.

2.5 Robustness Checks

This section presents a variety of robustness checks of our main analyses. The top row of Panel A of Table 2.8 (labelled "Base Case Multiple Predictors") corresponds to the bottom-right elements of Panels A, B, and C, of Table 2.3, respectively. The figures reported are the average adjusted R^2 from the multiple predictors regression and the number of coefficients (across the individual assets in a given asset class) that are significantly positive and significantly negative.

First, we consider varying the length of the window over which we compute return autocorrelations, which form the basis of our hedge fund illiquidity index. Our baseline analysis uses 12 months, and this table considers the use of 9 months, 18 months, and 24 months. The trade-off here is estimation error (longer windows have less estimation error) against timeliness (shorter windows are less stale). For international equities and currencies, we see that a 12-month window strikes the optimum balance between these competing goals, outperforming the other choices, particularly the shorter 9-month window. The results for U.S. corporate bonds, on the other hand, are reasonably insensitive to the choice of window length.

We next consider alternative models for obtaining an estimate of the degree of autocorrelation in hedge fund returns. Our baseline model uses simple first-order autocorrelation, which can be interpreted as the autoregressive parameter in an AR(1) model. We consider variations based on an AR(2), an MA(1), and an MA(2), all estimated using a rolling 12-month window. For the second-order models, there are two parameters that capture autocorrelation, and to summarize these into a single number, we use the R^2 from the model. For the MA(1) model, we simply use the estimated MA parameter. Table 2.8 shows that the second-order models tend

Table 2.8: Robustness Checks

This table presents robustness checks of our main results. The first row in Panel A corresponds to the bottom-right elements of Panels A, B, and C, of Table 2.3. We present the average adjusted R^2 of multiple predictor regressions across all assets within an asset class, the number of assets for which the coefficient of the hedge fund illiquidity index is significantly (at the 10% level) positive and significantly negative. All predictors are lagged by one month. The dependent variables are the monthly log excess returns, except for currencies, where we use the monthly log return. The time series for U.S. corporate bonds starts in January 1997 and ends in December 2011. For currencies and international equities, the time series starts in January 1995 and ends in December 2011. The time series for all hedge fund illiquidity indices start before January 1995, except for *AR(1) 18M*, it starts in June 1995, and for *AR(1) 24M*, it starts in December 1995. The style specific hedge fund illiquidity indices are computed only with funds of the corresponding investment style. Panel B reports the number of assets for which the hedge fund illiquidity index is significant at the 10% level in-sample when adding the 12 month autocorrelation of various risk factors as controls. The first number stands for positively significant coefficients and the second for negatively significant coefficients.

Panel A: Variations of the Hedge Fund Illiquidity Index			
	Int. Equities	**U.S. Corp. Bonds**	**Currencies**
Base Case Multiple Predictors	7.133 (21/0)	12.186 (32/0)	4.110 (5/0)
Varying Window Length			
9 Months	4.686 (12/0)	12.494 (33/0)	3.202 (1/0)
18 Months	5.276 (14/0)	10.921 (28/0)	3.606 (2/0)
24M Months	4.281 (8/0)	10.597 (32/0)	3.212 (0/0)
Varying Autocorrelation			
AR(2)	4.842 (17/0)	10.729 (32/0)	3.194 (1/0)
MA(1)	7.922 (21/0)	12.418 (30/0)	4.714 (5/0)
MA(2)	5.501 (20/0)	10.966 (31/0)	3.335 (1/0)
Different Styles			
Directional Traders	6.401 (21/0)		
Security Selection	6.420 (21/0)		
Fixed Income		12.280 (34/0)	
Global Macro			3.196 (2/0)
Other Variations			
AUM Weighted	5.577 (17/0)	10.123 (17/0)	4.033 (4/0)
Untrimmed	7.647 (21/0)	12.034 (32/0)	4.513 (5/0)
Panel B: Autocorrelation of Risk Factors as Controls			
Base Case Single Predictor	(21/0)	(31/0)	(6/0)
12M Autocorr. as Control			
HML U.S.	(21/0)	(30/0)	(6/0)
Mkt-Rf U.S.	(21/0)	(31/0)	(3/0)
Momentum U.S.	(21/0)	(30/0)	(6/0)
SMB U.S.	(21/0)	(31/0)	(6/0)
HML Global	(21/0)	(31/0)	(6/0)
Mkt-Rf Global	(21/0)	(33/0)	(6/0)
SMB Global	(21/0)	(31/0)	(6/0)
WML Global	(21/0)	(31/0)	(6/0)
PTFSBD	(21/0)	(28/0)	(6/0)
PTFSCOM	(21/0)	(31/0)	(3/0)
PTFSFX	(21/0)	(29/0)	(6/0)

to do worse than the first-order models, likely due to the increased estimation error from the additional parameter. The MA(1) model actually performs slightly better than our baseline AR(1) model, though the gains are small.

Third, we consider using hedge fund illiquidity indices based only on funds that are in style categories that are close to the asset class under consideration.[9] For international equities, we consider two hedge fund styles, namely Directional Traders and Security Selection. For corporate bonds, the natural style to consider is Fixed Income, and for currencies, we use Global Macro. For international equities, we find that the two style-specific illiquidity indices generate the same number of significant coefficients but a slightly lower average adjusted R^2. For U.S. corporate bonds, the Fixed Income hedge fund illiquidity index generates slightly stronger results than that based on all funds, while for currencies, the Global Macro index performs slightly worse. In Appendix 2.B, we present a theoretical analysis of the choice between asset-specific indices and a broader "all" index, highlighting the trade-offs between averaging across more funds and gaining robustness to style-misclassification, against greater precision of the illiquidity information.

Finally, we consider two further variations of our index. The first is based on estimated autocorrelations that are "untrimmed", in contrast to our baseline analysis which imposes that the autocorrelations are weakly positive. The second uses AUM weighting rather than equally weighting to construct the index. We see that the index based on untrimmed autocorrelation performs approximately as well as our baseline index, and in fact, slightly better for international equities and currencies. Constructing the index using AUM weights slightly worsens performance relative to the baseline case, as consistent with Figure 2.1, which shows that the AUM weighted index had smaller fluctuations than the equally weighted index. As noted above, this is not that surprising, as larger funds may be less sensitive to funding liquidity pressures imposed by prime-brokerage relationships and capital inflows.

[9]Details on the mapping from individual fund styles, as reported for a given hedge fund database, to the ten style categories that we consider in this chapter, are provided in Patton et al. (2013).

Panel B of Table 2.8 reports the number of significantly positive and significantly negative coefficients of the hedge fund illiquidity index in-sample, when including the 12-month autocorrelation of various risk factors as controls in the regression. The top row (labelled "Base Case Single Predictor") reports the results for the single predictor in-sample regressions. The performance of our hedge fund illiquidity index is robust to the inclusion of the risk factors. This indicates that the autocorrelation in hedge fund returns is not simply driven by the autocorrelation of risk factors.

2.6 Conclusion

Detecting evidence of hedge funds' impact on asset markets is an important endeavour given their size, leverage, and significant role in the provision of liquidity. We create a simple measure of aggregate hedge fund illiquidity by averaging fund-specific return autocorrelations across a large universe of hedge funds. We find that the resulting measure is a highly significant and robust predictor of returns, both in-sample and out-of-sample, for international equity indices, U.S. corporate bonds, and currencies.

We build a simple model of liquidity provision. A hedge fund, which provides liquidity to noise traders, is endowed with illiquid asset holdings and faces a shortfall penalty for not holding sufficient cash to cover the threat of redemptions by investors. The model is able to explain our main empirical findings, and yields additional testable implications which are supported in the data.

We view these results as a useful addition to the literature on hedge funds and the effects of asset-market liquidity on returns. In future work, we hope to explore the implications of this and other measures regarding the impact of hedge funds on asset markets to explain a broader range of asset market outcomes.

Appendix

2.A Extension of Clark and West (2006)

2.A.1 Properties of the extension

Clark and West (2006) consider the problem of testing equal predictive accuracy of a linear regression and a model with no parameters. We propose a simple generalisation that allows the "smaller" model to include a constant. Denote the two forecasts as

$$\hat{Y}_{t+1|t}^{(1)} = \hat{\gamma}_t \tag{2.A.1}$$

$$\hat{Y}_{t+1|t}^{(2)} = X_{t+1}'\hat{\beta}_t. \tag{2.A.2}$$

The second forecast is based on a set of predictor variables X_{t+1}, which includes a constant, and so nests the smaller model. Both forecasts are based on parameters estimated using a rolling window of data (and so they are not constant through the sample) of fixed length R. The sample period runs from $t = 1, 2, ..., R, R + 1, ..., R + P$. The mean squared prediction errors (MSPEs) over the out-of-sample period ($t = R + 1, ..., R + P$) are

$$\hat{\sigma}_1^2 = \frac{1}{P}\sum_{t=R+1}^{R+P}(y_{t+1} - \hat{\gamma}_t)^2 \tag{2.A.3}$$

$$\hat{\sigma}_2^2 = \frac{1}{P}\sum_{t=R+1}^{R+P}\left(y_{t+1} - X_{t+1}'\hat{\beta}_t\right)^2. \tag{2.A.4}$$

The null and alternative hypotheses are

$$H_0 : \sigma_1^2 = \sigma_2^2 \quad \text{vs.} \quad H_1 : \sigma_1^2 > \sigma_2^2. \tag{2.A.5}$$

That is, the small and large models are equally good under the null, while under the alternative, the small model is worse than the large model. Note that under the null, where the smaller model is correct, we have

$$y_{t+1} = \gamma^* + \epsilon_{t+1}, \text{ where } E_t\left[\epsilon_{t+1}\right] = 0. \tag{2.A.6}$$

The difference in the MSPEs of these two forecasts is

$$\begin{aligned}\hat{\sigma}_1^2 - \hat{\sigma}_2^2 &= \frac{1}{P}\sum_{t=R+1}^{R+P}\left\{(y_{t+1} - \hat{\gamma}_t)^2 - \left(y_{t+1} - X_{t+1}'\hat{\beta}_t\right)^2\right\} \\ &= \frac{1}{P}\sum_{t=R+1}^{R+P}\hat{\gamma}_t^2 - \frac{2}{P}\sum_{t=R+1}^{R+P} y_{t+1}\hat{\gamma}_t - \frac{1}{P}\sum_{t=R+1}^{R+P}\left(X_{t+1}'\hat{\beta}_t\right)^2 + \frac{2}{P}\sum_{t=R+1}^{R+P} y_{t+1}X_{t+1}'\hat{\beta}_t,\end{aligned} \tag{2.A.7}$$

and so

$$\hat{\sigma}_1^2 - \hat{\sigma}_2^2 \xrightarrow{p} E\left[\hat{\gamma}_t^2\right] - E\left[\left(X_{t+1}'\hat{\beta}_t\right)^2\right] - 2\gamma^* E\left[\left(\hat{\gamma}_t - X_{t+1}'\hat{\beta}_t\right)\right] \text{ under } H_0 \text{ as } P \to \infty, \tag{2.A.8}$$

using the fact that $E\left[\epsilon_{t+1}X_{t+1}'\hat{\beta}_t\right] = E\left[\epsilon_{t+1}\hat{\gamma}_t\right] = 0$.

When $\hat{\gamma}_t = 0\ \forall t$ and $\gamma^* = 0$, we are in the framework of Clark and West (2006), and they note that:

$$\hat{\sigma}_1^2 - \hat{\sigma}_2^2 \xrightarrow{p} -E\left[\left(X_{t+1}'\hat{\beta}_t\right)^2\right] \text{ under } H_0 \text{ as } P \to \infty. \tag{2.A.9}$$

Hence, when the smaller model is correct, the difference in MSPEs will be centered on a *negative value*, even though the two models are both correct (the larger model nests the smaller model, so it is also correct). This makes the test conservative (since we use a standard Normal as the asymptotic distribution of the test statistic,

which is centered on zero), and so it will be hard to reject in favor of the larger model when it is correct. Clark and West suggest adjusting $\hat{\sigma}_2^2$ so that the difference between the MSPEs is centered on zero under the null:

$$\hat{\sigma}_{2,adj}^2 = \hat{\sigma}_2^2 - \frac{1}{P}\sum_{t=R+1}^{R+P}\left(X_{t+1}'\hat{\beta}_t\right)^2. \tag{2.A.10}$$

They show that this provides better properties under the null and under the alternative.

We extend this adjustment to allow the smaller model to include a constant: this means $\gamma^* \neq 0$. Intuitively, allowing for a constant adds a small amount of variability to Forecast 1, making the extra penalty faced by Forecast 2 slightly lower in relative terms. This introduces some cross-product terms that need to be handled. Let us define an adjusted difference in MSPEs to correct for these terms:

$$\begin{aligned}\hat{\Delta}_{adj} =& \hat{\sigma}_1^2 - \hat{\sigma}_2^2 - \frac{1}{P}\sum_{t=R+1}^{R+P}\hat{\gamma}_t^2 + \frac{1}{P}\sum_{t=R+1}^{R+P}\left(X_{t+1}'\hat{\beta}_t\right)^2 \\ &+ 2\hat{\gamma}^*\frac{1}{P}\sum_{t=R+1}^{R+P}\left(\hat{\gamma}_t - X_{t+1}'\hat{\beta}_t\right),\end{aligned} \tag{2.A.11}$$

where

$$\hat{\gamma}^* = \frac{1}{P}\sum_{t=R+1}^{R+P} y_{t+1}. \tag{2.A.12}$$

Next we need to get a limiting distribution for $\hat{\Delta}_{adj}$. In the Clark-West case, $\hat{\Delta}_{adj}$ is just a linear combination of sample averages, and so this can be obtained by defining an "adjusted difference in per-period loss variable" ($\hat{f}_{t+1}$ in their equation 3.3), and conducting a t-test that that variable is zero mean. In our case, the cross-product term introduces a *product* of sample averages, and so we cannot use that approach. Instead, we use the delta method to get the limiting distribution of our

test statistic. Let

$$\underset{(6\times1)}{g_t} = \left[\left(y_{t+1} - \hat{\gamma}_t\right)^2, -\left(y_{t+1} - X'_{t+1}\hat{\beta}_t\right)^2, -\hat{\gamma}_t^2, \left(X'_{t+1}\hat{\beta}_t\right)^2, 2\left(\hat{\gamma}_t - X'_{t+1}\hat{\beta}_t\right), y_{t+1}\right]'. \tag{2.A.13}$$

Under standard regularity conditions we have

$$\sqrt{P}\left(\bar{g}_P - \bar{g}_0\right) \xrightarrow{d} N\left(0, V_g\right) \text{ as } P \to \infty \tag{2.A.14}$$

As in Clark and West (2006), the asymptotic covariance matrix, V_g, should be estimated using a HAC estimator (see, for example, Newey and West (1987)). Our test statistic is a nonlinear function of $\bar{g}_P$:

$$\hat{\Delta}_{adj} = f\left(\bar{g}_P\right) = \bar{g}_1 + \bar{g}_2 + \bar{g}_3 + \bar{g}_4 + \bar{g}_5\bar{g}_6, \tag{2.A.15}$$

and so by the delta method:

$$\sqrt{P}\left(f\left(\bar{g}_P\right) - f\left(g_0\right)\right) = \sqrt{P}\left(\hat{\Delta}_{adj} - \Delta_0\right) \xrightarrow{d} N\left(0, \nabla_g f\left(\bar{g}_P\right) V_g \nabla_g f\left(\bar{g}_P\right)'\right), \tag{2.A.16}$$

where $\nabla_g f\left(\bar{g}_P\right) = \partial f\left(\bar{g}_P\right)/\partial g' = [1, 1, 1, 1, \bar{g}_6, \bar{g}_5]$. We can use this to obtain a Clark-West style test for the larger model versus a model just including a constant. Specifically, we compute the test statistic:

$$\frac{\sqrt{P}\hat{\Delta}_{adj}}{\sqrt{\nabla_g f\left(\bar{g}_P\right) \hat{V}_g \nabla_g f\left(\bar{g}_P\right)'}}. \tag{2.A.17}$$

This test statistic is then compared to the right-tail critical values of the standard Normal distribution (1.28, 1.65, 2.33) to get a test at the 10%, 5%, or 1% level.

2.A.2 Monte Carlo simulation study

To study the finite-sample size of this extension of Clark and West (2006), we conduct a small simulation study. The target variable and forecasts are generated as follows:

$$\begin{aligned} Y_t &= \gamma^* + \beta^* \bar{X}_t + \epsilon_t, \\ \text{where } \bar{X}_t &\equiv \frac{1}{k}\sum_{i=1}^{k} X_{it}, \\ [X_t', \epsilon_t]' &\overset{i.i.d.}{\sim} N\left(0, I_{k+1}\right), \\ \hat{Y}^{(1)}_{t+1|t} &= \hat{\gamma}_t, \\ \text{and } \hat{Y}^{(2)}_{t+1|t} &= \left[1, X_{t+1}'\right]\hat{\beta}_t. \end{aligned} \tag{2.A.18}$$

We set $\beta^* = 0$, and so the DGP includes just a constant, γ^*. In this case both models are correct, but Forecast 1 will do better in finite samples because it does not include any irrelevant variables. To study power, we let $\beta^* = 0.1$, which means that each of the included regressors in Forecast 2 is useful, and Forecast 1 is misspecified.

We consider in-sample estimation period lengths of $R \in \{100, 200, 500, 1000, 5000\}$, and an out-of-sample evaluation period of $P = 1000$. The number of extra regressors in the larger model is set to $k \in \{1, 5, 10\}$. We repeat the simulation 1000 times.

In Table 2.9, we present the finite-sample rejection probabilities when the null hypothesis is true, using 0.05 level tests. We report our proposed "adjusted" test as well as the usual "unadjusted" test, which corresponds to the Diebold and Mariano (1995) test. Panel A shows that the unadjusted test is very conservative (with rejection probabilities much lower than 0.05), particularly for small R. The "adjusted" test provides better size control, with rejection probabilities closer to the nominal 0.05 level. The improvement from the adjustment is particularly great when k is large and R is small, consistent with the theory.

Panel B reports rejection probabilities when the null is false, which provides insights into the power of the adjusted and unadjusted tests. Consistent with the

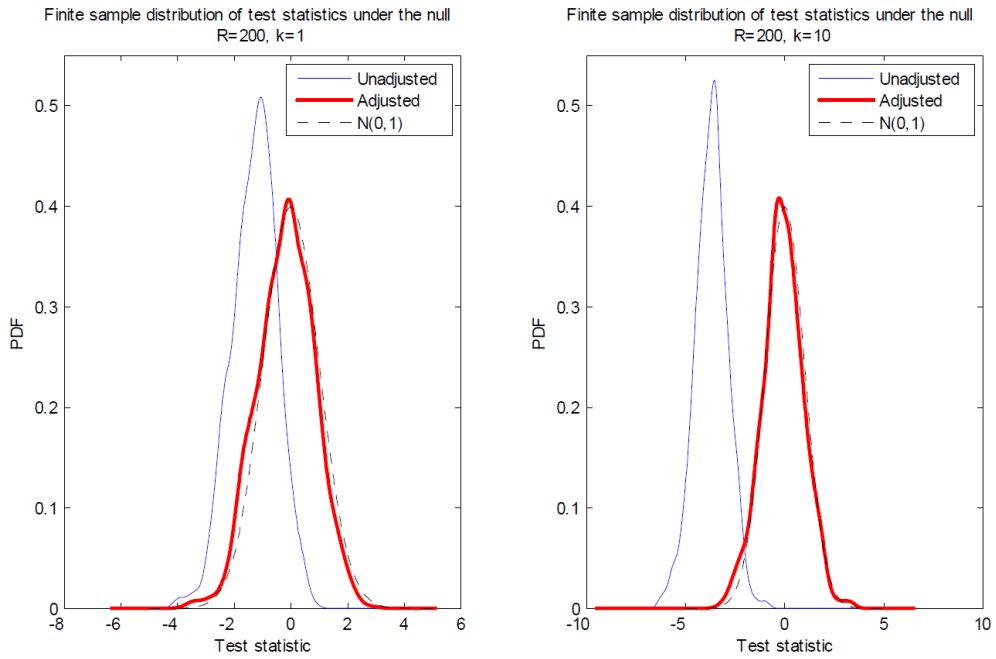
results under the null, the unadjusted test has low power to reject the null when it is false, and we see that the adjusted test has much better power than the unadjusted test. Figure 2.5 illustrates where the gains come from: the adjusted test re-centers the distribution of test statistics on zero, which makes the test less conservative under the null, and more powerful under the alternative.

Table 2.9: Rejection Probabilities for the Clark-West Extension

This table presents the rejection probabilities from the "unadjusted" and "adjusted" tests of equal predictive accuracy, at the 0.05 level. R represents the length of the in-sample estimation period, and k represents the number of regressors in the larger model (excluding the constant).

Panel A: True Null Hypothesis					
	R=100	R=200	R=500	R=1000	R=5000
			k=1		
Unadjusted	0.000	0.000	0.001	0.002	0.013
Adjusted	0.030	0.027	0.031	0.019	0.031
			k=5		
Unadjusted	0.000	0.000	0.000	0.000	0.013
Adjusted	0.049	0.035	0.042	0.050	0.049
			k=10		
Unadjusted	0.000	0.000	0.000	0.000	0.009
Adjusted	0.057	0.042	0.047	0.046	0.043
Panel B: False Null Hypothesis					
	R=100	R=200	R=500	R=1000	R=5000
			k=1		
Unadjusted	0.014	0.079	0.249	0.382	0.450
Adjusted	0.631	0.741	0.846	0.906	0.923
			k=5		
Unadjusted	0.014	0.364	0.869	0.949	0.962
Adjusted	0.997	1.000	1.000	1.000	1.000
			k=10		
Unadjusted	0.004	0.599	0.996	0.998	1.000
Adjusted	1.000	1.000	1.000	1.000	1.000

Figure 2.5: Test Statistic Distribution of the Clark-West Extension



2.B Hedge Fund Style Indices vs. an Aggregate Index

This section presents a stylized model to study how hedge fund style drift or style misreporting can affect the informativeness of an illiquidity index based on funds from a single style category relative to an aggregate index based on all hedge funds.

Consider the following simple framework. There are N individual hedge funds (with N large; in our empirical application it is 29,496). There are $S = 10$ different styles (to match our empirical application), and we will assume that there are $K = N/S$ funds in each style.

Assume that for each hedge fund style there is an asset index that is "close" to that style, so there are also S asset indices. For example, Fixed Income hedge funds and U.S. corporate bond returns or Security Selection hedge funds and equity returns. Returns on each asset class index are linked to liquidity, and have a component that is not related to liquidity:

$$r_{s,t+1} = \alpha_s^* + \beta_s^* L_{s,t}^* + u_{s,t+1}^*, \text{ where } s = 1, 2, ..., S. \tag{2.B.1}$$

We will assume that all parameters are the same across styles, all shocks are mean zero, Normal, homoskedastic, and independent from each other.

Hedge fund returns for funds in a given style provide us with a noisy estimate of the liquidity of that asset class:

$$\begin{aligned} &\tilde{L}_{i,t} = L_{s(i),t}^* + \eta_{i,t}, \text{ for } i = 1, 2, .., N \\ &\text{where } \ s(i) \in \{1, 2, ..., S\}, \end{aligned} \tag{2.B.2}$$

with $s(i)$ being the style of hedge fund i. In our empirical application, our noisy estimate of hedge fund illiquidity is based on autocorrelations; we abstract from that particular measure here and simply consider a generic noisy measure of liquidity.

We average the liquidity estimates from each fund in style s to obtain an aggregate liquidity index for that style:

$$\bar{L}_{s,t} = \frac{1}{K}\sum_{i=1}^{N} \tilde{L}_{i,t}\mathbf{1}\left\{s\left(i\right)=s\right\}, \tag{2.B.3}$$

and across all funds to get an aggregate liquidity index:

$$\bar{L}_{t} = \frac{1}{N}\sum_{i=1}^{N} \tilde{L}_{i,t}. \tag{2.B.4}$$

We use the (noisy) liquidity indices extracted from hedge fund returns in a regression to try to predict future asset index returns:

$$r_{s,t+1} = \tilde{\alpha}_s + \tilde{\beta}_s \bar{L}_{s,t} + \tilde{u}_{s,t+1}, \text{ where } s = 1, 2, ..., S. \tag{2.B.5}$$

Clearly, the closer $\bar{L}_{s,t}$ is to $L^*_{s,t}$, the better the prediction.

Case 1 (Styles are correctly reported): In this baseline case, we have

$$\begin{aligned}
\bar{L}_{s,t} &= \frac{1}{K}\sum_{i=1}^{N}\left(L^*_{s,t}+\eta_{i,t}\right)\mathbf{1}\left\{s\left(i\right)=s\right\} \\
&= L^*_{s,t} + \frac{1}{K}\sum_{i=1}^{N}\eta_{i,t}\mathbf{1}\left\{s\left(i\right)=s\right\} \\
&\equiv L^*_{s,t} + \epsilon_{s,t}, \text{ where } \epsilon_{s,t} \sim N\left(0, \frac{1}{K}\sigma^2_\eta\right).
\end{aligned} \tag{2.B.6}$$

The estimated coefficient on $\bar{L}_{s,t}$ is

$$\tilde{\beta}_s = \frac{Cov\left[\bar{L}_{s,t}, r_{s,t+1}\right]}{V\left[\bar{L}_{s,t}\right]} = \frac{Cov\left[L^*_{s,t}+\epsilon_{s,t}, r_{s,t+1}\right]}{V\left[L^*_{s,t}+\epsilon_{s,t}\right]} = \beta^*_s \frac{1}{1+\sigma^2_\eta/\left(K\sigma^2_L\right)}, \tag{2.B.7}$$

where $\sigma^2_L = V\left[L^*_{s,t}\right]$. This is the familiar shrinkage of a regression parameter towards

zero, when the dependent variable is measured with error. The R^2 of this model is

$$\tilde{R}^2 = Corr\left[\bar{L}_{s,t}, r_{s,t+1}\right]^2 = \frac{Cov\left[L^*_{s,t}, r_{s,t+1}\right]^2}{\left(\sigma_L^2 + \frac{1}{K}\sigma_\eta^2\right)\sigma_r^2} = R^{*2}\frac{1}{1+\sigma_\eta^2/\left(K\sigma_L^2\sigma_r^2\right)}, \tag{2.B.8}$$

where R^{*2} is the R^2 from the regression, if we could directly observe $L^*_{s,t}$, and $\sigma_r^2 = V\left[r_{s,t+1}\right]$. Next, we compare the above base case with what is obtained in the presence of style mis-labelling, style drift, etc.

Case 2 (Styles are not correctly reported): Perhaps because of style drift, or style mis-reporting, or just errors in classifying the hedge fund style, it may be that a fund that is listed as being in style s is actually in style j. Let p be the probability that the reported style label, denoted $\tilde{s}(i)$, is *correct*:

$$p = \Pr\left[\tilde{s}(i) = s(i)\right], \tag{2.B.9}$$

and assume this is the same across funds. For simplicity, assume that when the label is *incorrect*, the label is randomly chosen from the remaining $S-1$ styles.

To obtain the main results for this case, we will use the fact that

$$\begin{aligned}
\frac{1}{N}\sum_{i=1}^{N}\mathbf{1}\left\{\tilde{s}(i) = s, s(i) = s\right\} &\approx E\left[\mathbf{1}\left\{\tilde{s}(i) = s, s(i) = s\right\}\right] \quad \text{for } N \text{ large} \\
&= \Pr\left[\tilde{s}(i) = s, s(i) = s\right] \\
&= \Pr\left[\tilde{s}(i) = s | s(i) = s\right]\Pr\left[s(i) = s\right] \\
&= p\frac{1}{S}
\end{aligned} \tag{2.B.10}$$

and

$$\begin{aligned}\frac{1}{N}\sum_{i=1}^{N}\mathbf{1}\left\{\tilde{s}\left(i\right)=s,s\left(i\right)=j\neq s\right\} &\approx E\left[\mathbf{1}\left\{\tilde{s}\left(i\right)=s,s\left(i\right)=j\neq s\right\}\right]\\ &=\Pr\left[\tilde{s}\left(i\right)=s,s\left(i\right)=j\neq s\right]\\ &=\Pr\left[\tilde{s}\left(i\right)=s|s\left(i\right)=j\neq s\right]\Pr\left[s\left(i\right)=j\neq s\right]\\ &=\frac{1-p}{S-1}\frac{1}{S}.\end{aligned} \tag{2.B.11}$$

So we can obtain

$$\begin{aligned}\bar{L}_{s,t} &= \frac{1}{K}\sum_{i=1}^{N}\left(L^{*}_{s(i),t}+\eta_{i,t}\right)\mathbf{1}\left\{\tilde{s}\left(i\right)=s\right\}\\ &= \sum_{j=1}^{S}L^{*}_{j,t}\frac{1}{K}\sum_{i=1}^{N}\mathbf{1}\left\{\tilde{s}\left(i\right)=s,s\left(i\right)=j\right\}+\underbrace{\frac{1}{K}\sum_{i=1}^{N}\eta_{i,t}\mathbf{1}\left\{\tilde{s}\left(i\right)=s\right\}}_{\equiv\epsilon_{s,t}\sim N\left(0,\sigma_{\eta}^{2}/K\right)}\\ &= L^{*}_{s,t}\frac{1}{K}\sum_{i=1}^{N}\mathbf{1}\left\{\tilde{s}\left(i\right)=s,s\left(i\right)=s\right\}+\sum_{j=1,j\neq s}^{S}L^{*}_{j,t}\frac{1}{K}\sum_{i=1}^{N}\mathbf{1}\left\{\tilde{s}\left(i\right)=s,s\left(i\right)=j\right\}+\epsilon_{s,t}\\ &= L^{*}_{s,t}\frac{N}{K}\frac{p}{S}+\sum_{j=1,j\neq s}^{S}L^{*}_{j,t}\frac{N}{K}\frac{1-p}{S-1}\frac{1}{S}+\epsilon_{s,t}\\ &= pL^{*}_{s,t}+\frac{1-p}{S-1}\sum_{j=1,j\neq s}^{S}L^{*}_{j,t}+\epsilon_{s,t}\end{aligned} \tag{2.B.12}$$

using the fact that $N/K=S$. The last line expresses $\bar{L}_{s,t}$ as a weighted average of each of the individual style indices. We can rewrite it to be a weighted average of

just the target style index and the "all" index:

$$\begin{aligned}
\bar{L}_{s,t} &= pL^*_{s,t} + \frac{1-p}{S-1}\sum_{j=1,j\neq s}^{S} L^*_{j,t} + \epsilon_{s,t} \\
&= pL^*_{s,t} + \frac{1-p}{S-1}\sum_{j=1}^{S} L^*_{j,t} + \epsilon_{s,t} - \frac{1-p}{S-1}L^*_{s,t} \\
&= \frac{pS-1}{S-1}L^*_{s,t} + \frac{1-p}{S-1}\sum_{j=1}^{S} L^*_{j,t} + \epsilon_{s,t} \\
&= \frac{pS-1}{S-1}L^*_{s,t} + \frac{S\,(1-p)}{S-1}\bar{L}^*_t + \epsilon_{s,t}.
\end{aligned} \tag{2.B.13}$$

Note that the noise term, $\epsilon_{s,t}$, does not depend on p, but it does depend on K through its variance. When $p = 1$, the style label is always correct and we obtain

$$\bar{L}_{s,t} = L^*_{s,t} + \epsilon_{s,t}, \tag{2.B.14}$$

as in the base case. When $p = 1/S$ the style label is correct as often as a random guess and we obtain

$$\bar{L}_{s,t} = \frac{\frac{1}{S}S-1}{S-1}L^*_{s,t} + \frac{S\left(1-\frac{1}{S}\right)}{S-1}\bar{L}^*_t + \epsilon_{s,t} = \bar{L}^*_t + \epsilon_{s,t}. \tag{2.B.15}$$

Hence, when the labels are randomly applied, the index average is equal to the aggregate index average plus some measurement error.

The above calculations reveal the "bias-variance" trade-off between using a style-specific index and using an aggregate index: when we average across all funds we get an index centered on the average liquidity, $\bar{L}^*_t$, which is not really what we want (a form of bias), but it has small measurement error (lower variance): $\epsilon_t \sim N\left(0, \sigma^2_\eta/N\right)$. When we average across just those funds in the style we care about, we get an index that is more heavily weighted on the style we care about (for values of $p > 1/S$), but it has greater measurement error: $\epsilon_{s,t} \sim N\left(0, \sigma^2_\eta/K\right)$.

To find the threshold value for p that makes these two approaches equally accu-

rate, we need to take a stand on the correlation between liquidity factors in different styles. To generate correlation between liquidity in different asset classes, we assume a simple factor structure. Let

$$L^*_{s,t} = \gamma L^*_{c,t} + \nu_{s,t}, \text{ where } s = 1, 2, ..., L. \tag{2.B.16}$$

Without loss of generality, let $V\left[L^*_{c,t}\right] = 1$, and then note that $\sigma^2_\nu = \sigma^2_L - \gamma^2$. Note that the R^2 of this regression is

$$\rho^2_c \equiv Corr\left[L^*_{s,t}, L^*_{c,t}\right]^2 = \frac{\gamma^2}{\gamma^2 + \sigma^2_v} = \frac{\gamma^2}{\sigma^2_L}. \tag{2.B.17}$$

Then we obtain

$$\begin{aligned}
Cov\left[L^*_{s,t}, L^*_{j,t}\right] &= \gamma^2, \;\; Cov\left[r_{s,t+1}, L^*_{s,t}\right] = \beta^* \sigma^2_L, \text{ and} \\
Cov\left[r_{s,t+1}, L^*_{j,t} | j \neq s\right] &= Cov\left[r_{s,t+1}, \gamma L^*_{c,t} + \nu_{j,t}\right] \\
&= Cov\left[r_{s,t+1}, \gamma\left(\gamma \frac{1}{\sigma^2_L} L^*_{s,t} + z_t\right) + \nu_{j,t}\right] \\
&= \gamma^2 \frac{1}{\sigma^2_L} Cov\left[r_{s,t+1}, L^*_{s,t}\right] \\
&= \rho^2_c \beta^* \sigma^2_L.
\end{aligned} \tag{2.B.18}$$

Thus, if the common component is very strong (ρ_c is near ± 1), then using the "wrong" style liquidity does not matter – it is almost as good as the right style. When the common component is weak, the covariance is pulled towards zero.

Now consider the R^2 from the predictive regression using the aggregate index. First, we obtain the covariance between the aggregate index and a given asset style return:

$$\begin{aligned}
Cov\left[\bar{L}_t, r_{s,t+1}\right] &= \frac{1}{S} Cov\left[L^*_{s,t}, r_{s,t+1}\right] + \frac{1}{S} \sum_{j=1, j\neq s}^{S} Cov\left[L^*_{j,t}, r_{s,t+1}\right] \\
&= \frac{1 + (S-1)\rho^2_c}{S} \beta^* \sigma^2_L.
\end{aligned} \tag{2.B.19}$$

and the variance of the aggregate index:

$$\begin{aligned} V\left[\bar{L}_t\right] &= \frac{1}{S^2} V\left[\sum_{j=1}^{S} L_{j,t}\right] + V\left[\epsilon_t\right] = \frac{1}{S^2} V\left[S\gamma L^*_{c,t}\right] + \frac{1}{S^2}\sum_{j=1}^{S} V\left[\nu_{j,t}\right] + \frac{1}{N}\sigma_\eta^2 \\ &= \frac{(S-1)K\gamma^2 + K\sigma_L^2 + \sigma_\eta^2}{N} \\ &= \frac{(S-1)}{S}\gamma^2 + \frac{1}{S}\sigma_L^2 + \frac{1}{N}\sigma_\eta^2. \end{aligned} \tag{2.B.20}$$

Combing the above, we find that the R^2 is

$$\begin{aligned} R^2_{all} &= \frac{\left(1+(S-1)\rho_c^2\right)^2 K\beta^{*2}\sigma_L^4}{(S-1)N\rho_c^2\sigma_L^2 + N\sigma_L^2 + S\sigma_\eta^2}\frac{1}{\sigma_r^2} \\ &= \left(1+(S-1)\rho_c^2\right)^2 \frac{K}{S}\frac{1}{\left((S-1)K\rho_c^2 + K + \lambda^2\right)} R^{*2}, \end{aligned} \tag{2.B.21}$$

where $\lambda^2 \equiv \frac{\sigma_\eta^2}{\sigma_L^2}$ and $R^{*2} \equiv Corr\left[L^*_{s,t}, r_{s,t+1}\right]^2 = \frac{\beta^{*2}\sigma_L^2}{\sigma_r^2}$.

We can contrast this R^2 with that obtained from the individual style liquidity index:

$$R_s^2 = \frac{p^2(S-1)K}{\left(p^2S - 2p + 1\right)K + (S-1)\lambda^2} R_s^{*2}. \tag{2.B.22}$$

Using these two expressions, we can solve for the value p^* that equates them. The solution is lengthy, and we illustrate it below. The comparative statics reveal that the required proportion of correct individual style labels (p^*) is higher when:

1. Holding the number of individual funds (N) fixed, the number of hedge funds in each style (K) is lower, or the number of styles (S) is higher: fewer funds per style means the individual style index is noisier.

2. The noise to signal ratio of each individual liquidity estimate ($\lambda \equiv \sigma_\eta/\sigma_L$) is higher.

3. The correlation between the individual style liquidity measures and the common liquidity measure (ρ_c) is higher. When this is higher, the aggregate index does better, and so the style labels must be accurate with higher probability

to counteract this.

In our data, the average correlation between our style indices is $\rho_L = 0.5869$, suggesting that $\rho_c = 0.7661$. Depending on the parameters that describe the model, it is possible that the individual style illiquidity index is always preferred, is never preferred, or is sometimes preferred to the aggregate index. Thus, which illiquidity index to use in practice is something that needs to be determined empirically.

2.C Numerical Procedure to Simulate Model

In our model, the hedge fund has mean-variance preferences, which take the following form for $\Phi_{max} \leq C_0 - x_1 p_1$:

$$x_1(E[d_2] - p_1) - \alpha \frac{1}{2} x_1^2 \sigma_d^2. \tag{2.C.1}$$

For the case where $\Phi_{max} > C_0 - x_1 p_1$, a shortfall penalty for not holding enough cash is added:

$$x_1(E[d_2] - p_1) - \alpha \frac{1}{2} x_1^2 \sigma_d^2 - \lambda \frac{1}{2} (\Phi_{max} - (C_0 - x_1 p_1))^2, \tag{2.C.2}$$

where $\lambda > 0$. The demand function for the first case is

$$x_1 = \frac{E[d_2] - p_1}{\alpha \sigma_d^2}. \tag{2.C.3}$$

For the second case, the demand function is given by

$$x_1 = \frac{E[d_2] - p_1 - \lambda(\Phi_{max} - C_0) p_1}{\alpha \sigma_d^2 + \lambda p_1^2}. \tag{2.C.4}$$

Which demand function for x_1 is applicable depends again on x_1. Hence, we need a numerical procedure to simulate the model. The procedure has three steps:

1. Use the demand function $x_1 = \frac{E[d_2]-p_1}{\alpha\sigma_d^2}$, and determine p_1^* at which the market clears.

2. Check whether $\Phi_{max} \leq C_0 - x_1 p_1^*$. If the inequality holds, the solution is found. If the inequality does not hold, proceed with 3.

3. Use the demand function $x_1 = \frac{(E[d_2]-p_1)-\lambda(\Phi_{max}-C_0)p_1}{\alpha\sigma_d^2+\lambda p_1^2}$, and determine $\hat{p}_1$ at which the market clears. Based on Proposition 1, $\Phi_{max} > C_0 - x_1\hat{p}_1$ will hold conditional on $\Phi_{max} > C_0 - x_1 p_1^*$.

Proposition 1. *If $\hat{p}_1$ exists and $\lambda > 0$, $\Phi_{max} > C_0 - x_1 p_1^* \Rightarrow \Phi_{max} > C_0 - x_1\hat{p}_1$.*

Proof of Proposition 1. When the noise trader demand is negative, i.e. $u_1 < 0$ and $x_1 > 0$, for $\Phi_{max} > C_0 - x_1\hat{p}_1$ to hold, we need

$$\frac{\Phi_{max} - C_0}{-x_1} < \hat{p}_1. \tag{2.C.5}$$

Using the quadratic formula for $\hat{p}_1$ given in equation (2.12), we can rewrite the inequality as

$$\begin{aligned}\hat{p}_1 &= \frac{1+\lambda(\Phi_{max}-C_0) - \sqrt{(1+\lambda(\Phi_{max}-C_0))^2 - 4(\lambda u_1)(E[d_2]+\alpha\sigma_d^2 u_1)}}{2\lambda u_1} \\ &> \frac{\Phi_{max}-C_0}{-x_1}.\end{aligned} \tag{2.C.6}$$

Substituting $-u_1$ for x_1 and multiplying both sides by u_1 (note that $u_1 < 0$) results in

$$\frac{1+\lambda(\Phi_{max}-C_0) - \sqrt{(1+\lambda(\Phi_{max}-C_0))^2 - 4(\lambda u_1)(E[d_2]+\alpha\sigma_d^2 u_1)}}{2\lambda} < \Phi_{max} - C_0. \tag{2.C.7}$$

The inequality can be rewritten as

$$1 - \lambda(\Phi_{max}-C_0) - \sqrt{(1+\lambda(\Phi_{max}-C_0))^2 - 4(\lambda u_1)(E[d_2]+\alpha\sigma_d^2 u_1)} < 0. \tag{2.C.8}$$

We know that $p_1^* = E[d_2] + \alpha\sigma_d^2 u_1$ from equation (2.10). Assuming that $p_1^* = \frac{\Phi_{max}-C_0}{-x_1} = E[d_2] + \alpha\sigma_d^2 u_1$ and using $x_1 = -u_1$:

$$\begin{aligned}
&1 - \lambda(\Phi_{max} - C_0) - \sqrt{(1 + \lambda(\Phi_{max} - C_0))^2 + 4(\lambda x_1)(\frac{\Phi_{max} - C_0}{-x_1})} \\
&= 1 - \lambda(\Phi_{max} - C_0) - \sqrt{(1 + \lambda(\Phi_{max} - C_0))^2 - 4(\lambda)(\Phi_{max} - C_0)} \\
&= 1 - \lambda(\Phi_{max} - C_0) - \sqrt{(1 - \lambda(\Phi_{max} - C_0))^2} \\
&= 0.
\end{aligned} \tag{2.C.9}$$

Hence, for $\Phi_{max} > C_0 - x_1 p_1^*$, which implies $\frac{\Phi_{max}-C_0}{-x_1} < p_1^*$, and $\lambda > 0$, the inequality

$$1 - \lambda(\Phi_{max} - C_0) - \sqrt{(1 + \lambda(\Phi_{max} - C_0))^2 + 4(\lambda x_1)p_1^*} < 0 \tag{2.C.10}$$

will hold.

When the noise trader demand is positive, i.e. $u_1 > 0$ and $x_1 < 0$, for $\Phi_{max} > C_0 - x_1\hat{p}_1$ to hold, we need

$$\frac{\Phi_{max} - C_0}{-x_1} > \hat{p}_1. \tag{2.C.11}$$

Via the same transformation as above, we have the following inequality which ensures that $\Phi_{max} > C_0 - x_1\hat{p}_1$:

$$1 - \lambda(\Phi_{max} - C_0) - \sqrt{(1 + \lambda(\Phi_{max} - C_0))^2 + 4(\lambda x_1)(E[d_2] + \alpha\sigma_d^2 u_1)} < 0. \tag{2.C.12}$$

For $\frac{\Phi_{max}-C_0}{-x_1} = p_1^*$, the term on the left hand side will equal 0, as derived in equation (2.C.9). As $x_1 < 0$ and $\lambda > 0$, $p_1^* < \frac{\Phi_{max}-C_0}{-x_1}$ will make the inequality hold. □

2.D Additional Tables

Table 2.10: In-Sample Predictive Performance for U.S. Corporate Bonds

Reported are the signs of the predictors (in Panel A, only the sign of the hedge fund illiquidity index), the adjusted R^2 of the in-sample predictive regression in %, and the significance of the hedge fund illiquidity index (* for 10% and ** for 5%). AUM weighted flows are not included in Panel A. The time series starts in January 1997 and ends in December 2011. Newey-West standard errors are used.

Panel A: Multiple Predictors		
Portfolios	**Excl. Hedge Fund Illiq. Index**	**Incl. Hedge Fund Illiq. Index**
AAA 1-3Y	3.997	(+) 3.451
AAA 3-5Y	4.055	(+) 3.659
AAA 5-7Y	4.219	(+) 3.721
AAA 7-10Y	2.924	(+) 2.550
AAA 10-15Y	3.624	(+) 3.202
AAA 15+Y	3.114	(+) 2.766
AA 1-3Y	2.359	(+) 3.202
AA 3-5Y	2.586	(+) 3.735
AA 5-7Y	0.996	(+) 2.164
AA 7-10Y	3.177	(+) 4.590*
AA 10-15Y	0.191	(+) 0.252
AA 15+Y	2.628	(+) 4.426**
A 1-3Y	8.170	(+) 10.958**
A 3-5Y	3.312	(+) 6.046**
A 5-7Y	3.131	(+) 6.741**
A 7-10Y	3.890	(+) 6.727**
A 10-15Y	0.868	(+) 3.661**
A 15+Y	3.288	(+) 6.446**
BBB 1-3Y	12.754	(+) 19.073**
BBB 3-5Y	10.003	(+) 15.307**
BBB 5-7Y	9.147	(+) 14.719**
BBB 7-10Y	5.899	(+) 11.584**
BBB 10-15Y	3.796	(+) 9.129**
BBB 15+Y	4.486	(+) 10.086**
BB 1-3Y	13.833	(+) 17.997**
BB 3-5Y	14.102	(+) 19.380**
BB 5-7Y	13.710	(+) 19.447**
BB 7-10Y	11.395	(+) 16.892**
BB 10-15Y	11.108	(+) 14.353**
BB 15+Y	17.380	(+) 21.777**
B 1-3Y	11.778	(+) 17.050**
B 3-5Y	21.044	(+) 25.874**
B 5-7Y	16.845	(+) 21.254**
B 7-10Y	15.466	(+) 20.282**
B 10-15Y	9.849	(+) 17.244**
B 15+Y	8.350	(+) 16.628**
C 1-3Y	8.609	(+) 11.306**
C 3-5Y	17.014	(+) 24.016**
C 5-7Y	24.610	(+) 30.098**
C 7-10Y	14.209	(+) 19.870**
C 10-15Y	20.998	(+) 22.167**
C 15+Y	12.422	(+) 17.989**

Panel B: Single Predictor

Portfolios	Hedge Fund Illiq. Index	Lagged Returns	Pastor-Stamb. Traded Liq. Factor	VIX AR(2) Shocks	VWM Excess Return	Hedge Fund Flows
AAA 1-3Y	(+) -0.56	(+) 1.14*	(-) -0.11	(+) 1.90**	(-) 4.83**	(-) -0.32
AAA 3-5Y	(+) -0.44	(+) 0.14	(-) 0.34	(+) 1.89**	(-) 4.97**	(-) -0.56
AAA 5-7Y	(+) -0.55	(+) -0.34	(-) -0.26	(+) 1.91**	(-) 5.58**	(+) -0.48
AAA 7-10Y	(+) -0.39	(+) -0.23	(-) 0.34	(+) 1.84**	(-) 3.99**	(+) -0.41
AAA 10-15Y	(+) -0.47	(+) -0.56	(-) 0.24	(+) 2.13*	(-) 4.68**	(+) -0.14
AAA 15+Y	(+) -0.27	(+) -0.44	(-) 1.89*	(+) 2.04*	(-) 2.87**	(+) 0.81
AA 1-3Y	(+) 1.33	(+) 1.46**	(-) 1.38	(+) -0.21	(-) 0.39	(-) 1.00
AA 3-5Y	(+) 1.55	(+) 0.65*	(-) 1.64*	(+) 0.01	(-) 1.22	(-) -0.31
AA 5-7Y	(+) 1.30	(+) -0.10	(-) 0.75	(+) -0.30	(-) 0.83	(-) -0.56
AA 7-10Y	(+) 1.88*	(+) 0.53	(-) 1.61*	(+) 0.59	(-) 1.96**	(+) -0.56
AA 10-15Y	(+) -0.05	(-) -0.45	(-) -0.10	(+) -0.11	(-) 1.08*	(+) 0.01
AA 15+Y	(+) 2.00**	(-) -0.56	(-) 3.41**	(+) -0.12	(-) 0.49	(-) -0.55
A 1-3Y	(+) 4.96**	(+) 7.65**	(-) 0.96	(-) -0.19	(+) -0.51	(-) 0.70
A 3-5Y	(+) 3.77**	(+) 2.53**	(-) 0.95	(-) -0.46	(-) -0.49	(-) 0.15
A 5-7Y	(+) 4.93**	(+) 2.39*	(-) 0.97	(-) -0.56	(-) -0.28	(-) 0.06
A 7-10Y	(+) 4.03**	(+) 2.17*	(-) 1.61*	(-) -0.51	(-) -0.37	(-) -0.40
A 10-15Y	(+) 3.08**	(+) 0.06	(-) 0.36	(-) -0.51	(-) -0.20	(-) -0.10
A 15+Y	(+) 4.09**	(+) 0.42	(-) 2.87**	(-) -0.46	(-) -0.47	(-) -0.38
BBB 1-3Y	(+) 9.23**	(+) 9.05**	(-) 0.23	(-) 4.75*	(+) 1.77	(-) 3.70*
BBB 3-5Y	(+) 7.31**	(+) 6.28**	(-) 0.02	(-) 3.36	(+) 0.41	(-) 1.95
BBB 5-7Y	(+) 7.73**	(+) 5.98**	(-) 0.16	(-) 3.02	(+) 0.47	(-) 0.99
BBB 7-10Y	(+) 7.02**	(+) 3.18*	(-) 0.22	(-) 2.30	(+) -0.01	(-) 0.84
BBB 10-15Y	(+) 5.92**	(+) 1.43*	(-) 0.05	(-) 1.39	(+) -0.47	(-) 0.59
BBB 15+Y	(+) 6.94**	(+) 1.21	(-) 1.91**	(-) 0.95	(+) -0.30	(-) 0.12

Panel B (Continued): Single Predictor						
Portfolios	**Hedge Fund Illiq. Index**	**Lagged Returns**	**Pastor-Stamb. Traded Liq. Factor**	**VIX AR(2) Shocks**	**VWM Excess Return**	**Hedge Fund Flows**
BB 1-3Y	(+) 5.95**	(+) 10.15**	(-) -0.14*	(-) 5.76*	(+) 2.42**	(-) 2.16*
BB 3-5Y	(+) 6.68**	(+) 8.53**	(-) 0.02	(-) 8.05*	(+) 3.44**	(-) 0.90
BB 5-7Y	(+) 7.32**	(+) 8.51**	(-) -0.29	(-) 8.12*	(+) 2.94**	(+) 1.22
BB 7-10Y	(+) 7.55**	(+) 6.91**	(-) 0.13	(-) 5.61	(+) 2.47*	(-) 1.56*
BB 10-15Y	(+) 3.85**	(+) 4.25**	(-) -0.16	(-) 9.18*	(+) 4.54**	(-) -0.19
BB 15+Y	(+) 7.19**	(+) 12.33**	(-) 0.16	(-) 8.82**	(+) 6.63**	(-) 0.08
B 1-3Y	(+) 3.59**	(-) -0.35	(-) -0.54	(-) 9.60**	(+) 5.28**	(-) 1.24
B 3-5Y	(+) 6.97**	(+) 14.59**	(-) -0.40	(-) 14.64**	(+) 12.56**	(-) 0.27
B 5-7Y	(+) 5.02**	(+) 6.39**	(-) 0.03	(-) 14.88**	(+) 10.13**	(-) -0.49
B 7-10Y	(+) 6.46**	(+) 7.05**	(-) 0.64	(-) 10.85*	(+) 7.34**	(-) 0.32*
B 10-15Y	(+) 8.48**	(+) 1.79	(-) 0.92	(-) 7.49**	(+) 5.39**	(-) 0.61
B 15+Y	(+) 9.11**	(+) 1.48	(-) 0.34	(-) 6.72**	(+) 6.30**	(-) 1.07
C 1-3Y	(+) 2.70**	(+) 0.34	(-) -0.08	(-) 8.35*	(+) 4.48**	(-) -0.51
C 3-5Y	(+) 9.12**	(+) 12.40**	(-) -0.51	(-) 12.51**	(+) 12.14**	(-) -0.32
C 5-7Y	(+) 8.30**	(+) 16.77**	(-) -0.19	(-) 16.13**	(+) 14.79**	(+) -0.39
C 7-10Y	(+) 7.84**	(+) 8.23**	(-) 0.42	(-) 8.49**	(+) 10.89**	(-) 0.10
C 10-15Y	(+) 1.89	(+) 18.35**	(+) -0.08	(-) 9.42**	(+) 10.50**	(-) 1.51
C 15+Y	(+) 6.55**	(+) 6.68**	(-) -0.50	(-) 10.73**	(+) 8.58**	(-) -0.37

Chapter 3

On the Time-Varying Risk of Hedge Fund Portfolios and Risk-Shifting

Abstract

Measuring the risk of hedge fund portfolios is important for regulators and investors alike. Since the information on hedge fund holdings is very restricted, researchers have used the variance of returns as a proxy for risk. However, estimating the time-varying variance of hedge fund returns is a challenge, as for most hedge funds only a short time series of monthly data is available. In this chapter, I propose a new method for measuring the time-varying variance, which is based on recent advances in the conditional variance and correlation estimation literature. I use this method to investigate whether equity long-short hedge funds engage in risk-shifting driven by their past performance relative to their peers. I find that hedge funds which have strongly underperformed or outperformed their peers in recent months increase their exposure to the core strategy, i.e. the equity long-short strategy, and to non-core strategies. I also discover strong evidence that the risk shifting is mitigated for hedge funds with long redemption periods.

3.1 Introduction

Regulators and policymakers have had a strong interest in the hedge fund industry ever since the Long-Term Capital Management crisis of 1998. In recent years, their ambition to bring greater transparency to the hedge fund industry has led to enhanced regulation by the U.S. Securities and Exchange Commission (SEC) through the Dodd-Frank Wall Street Reform and Consumer Protection Act. The recurring collapses of large hedge funds indicate that a substantial proportion pursue high risk investment strategies. What drives the risk taking of hedge funds is an important question, as hedge funds play a central role in the global capital markets by employing parts of their aggregate U.S.$ 1.5 trillion of assets under management (AUM) to liquidity provision, as described in Chapter 2. Further, investors could benefit from insights into the risk taking behaviour of hedge fund managers to achieve a more effective allocation of assets across hedge funds. However, the research showing the drivers of risk in hedge fund portfolios is limited.

In this chapter, I analyse whether performance benchmarking – comparing the performance of a hedge fund relative to a benchmark – has an effect on subsequent risk taking. Is a hedge fund which underperforms its peers more likely to undertake riskier investments in the future to catch up? There is an extensive literature on mutual fund benchmarking (see, for example, Brown, Harlow, and Starks (1996) and Basak, Pavlova, and Shapiro (2007)), where the general consensus is that mutual funds which underperform in the first half of the year tend to increase the risk of their investments in the second half. The motivation to rank well relative to peers at the end of the year is due to the strong link between the end of year performance and future fund flows. As a measure of risk-shifting, the standard deviations of the mutual fund's return over the first and second halves of the year are compared. For hedge funds, similar studies have been conducted. Brown, Goetzmann, and Park (2001) find an association between past performance and risk levels in line with the mutual fund industry. Aragon and Nanda (2012) also conclude that risk-shifting

behaviour exists in the hedge fund industry, but that certain fund characteristics mitigate the phenomenon. A high water mark (HWM), a large personal stake, and a low likelihood of fund closure weaken the tendency of underperforming funds to increase risk. The empirical literature is complemented by theoretical work. For example, Hodder and Jackwerth (2007) build a theoretical model which explains the risk taking of hedge fund managers through an incentive contract.

A key difference between hedge fund data and mutual fund data is the frequency. The majority of databases provide hedge fund return data at a monthly frequency. Hence, a comparison between the return standard deviations of the first and the second halves of the year is based on a small number of data points and, consequently, noisy estimates. I propose a new method for measuring the time-varying risk of hedge funds. The method is rooted in the conditional variance and correlation estimation literature: it combines a panel GARCH model developed by Pakel, Shephard, and Sheppard (2011) with the dynamic conditional β introduced by Engle (2012). I focus on equity long-short hedge funds and measure each fund's time-varying return variance. In line with the existing literature, the return variance acts as a proxy for risk. Further, I decompose the variance into three components. The first component accounts for the variance of the equity long-short strategy. The second component represents the hedge fund's exposure to the equity long-short strategy, i.e. the core strategy. The third component is the variance which is due to investment strategies other than the core strategy, i.e. the non-core strategies.

The estimates of the time-varying risk are then used to investigate whether the performance relative to a benchmark affects the future risk taking of hedge funds. I find that funds which are far behind their peers and funds which are in front, on a year-to-date basis, tend to increase their exposure to the core strategy and non-core strategies over the subsequent months. The further the distance to the benchmark, in both directions, the more pronounced is the risk-shifting. The finding that underperforming hedge funds tend to increase the risk of their portfolios is in

line with Brown et al. (2001) and Aragon and Nanda (2012). However, evidence that hedge funds which are outperforming their peers tend to increase their risk exposure is new. Also, my results show that hedge funds with long redemption periods engage in less risk taking due to benchmarking. Interestingly, the HWM does not appear to act as a benchmark on the basis of which hedge fund managers adjust their risk.

The remainder of this chapter starts with Section 3.2, which describes the data sample. Then I outline the methodology in Section 3.3. Section 3.4 presents the results. Various robustness checks are reported in Section 3.5, and Section 3.6 concludes.

3.2 Data

3.2.1 Hedge fund returns

I start off with a panel data set of monthly hedge fund returns over the period from January 1994 to December 2012. This data set has been consolidated from the TASS, HFR, CISDM, Morningstar, and BarclayHedge databases. These data represent the most comprehensive set of hedge fund data available from public sources and comprise a total of 32,519 individual hedge funds, including all births and deaths of funds over the period. The details of the process followed to consolidate these data can be found in Patton, Ramadorai, and Streatfield (2013). Reported returns are net of management and incentive fees. The data set is very similar to that used in Chapter 2 of this thesis – the main difference is that I have added an additional year to the sample and discarded all of the hedge funds with less than 50 consecutive monthly returns, thereby eliminating roughly 50% of the them. The lower bound for the number of time series observations is motivated by the variance estimation methodology described in Section 3.3.

Getmansky, Lo, and Makarov (2004) find that the reported hedge fund returns

are often subject to return smoothing. They propose estimating the unsmoothed returns via a moving average model. I unsmooth all hedge fund returns in my data set using their proposed method: estimating the MA(2) coefficients on the de-meaned observed hedge fund returns and assuming that the observed return in t is an average of the true returns in t, $t-1$, and $t-2$. In Chapter 2, we did not unsmooth the returns, as the extent of the smoothing provided us with information regarding the illiquidity of the hedge fund portfolios. However, in this chapter, I am interested in the true returns, as the variance of the true returns is a better proxy for portfolio risk.

In my analysis, I focus on hedge funds with an equity long-short strategy. Fung and Hsieh (2004) demonstrate that equity long-short funds have a positive exposure to the two asset-based style factors S&P 500 returns and small cap minus large cap returns, SC-LC. To select the hedge funds which have implemented an equity long-short strategy over the sample period, I regress the returns of each hedge fund on the asset-based style factors of Fung and Hsieh (2004). The data are from David Hsieh's website. The funds which load positively and significantly at the 1% level on the S&P 500 returns and the SC-LC returns are selected. This leads to a subsample of 1,205 funds. Summary statistics for the reported and unsmoothed returns are provided in Panel A of Table 3.1.

3.2.2 Equity long-short strategy returns

To measure the exposure of equity long-short hedge funds to their core strategy over time, a return series for the strategy is required. For this purpose, I compute the equally weighted average of all 1,205 selected hedge fund returns in the sample. This equally weighted average is then regressed on a constant and the two asset-based style factors S&P 500 and the SC-LC, and the fitted returns are assumed to be the returns of an equity long-short strategy. Coefficient estimates and test statistics of the regression as well as summary statistics of the fitted returns are given in Panels

Table 3.1: Summary Statistics

Panel A reports the summary statistics for the reported and unsmoothed monthly returns averaged across 1,205 equity long-short hedge funds from January 1994 to December 2012. Panel B reports the parameter estimates from regressing the equally weighted average unsmoothed returns across the 1,205 hedge funds on a constant and the asset-based style factors S&P 500 and SC-LC. Newey-West standard errors are used. Panel C contains the summary statistics of the equity long-short strategy returns from January 1994 to December 2012. The strategy returns are the fitted values of the regression shown in Panel B. Significant coefficients at the 10% and 5% level are indicated by * and **, respectively.

Panel A: Equity Long-Short Hedge Fund Monthly Returns

	Averaged Across 1,205 Hedge Funds			
	Average	**Std. Dev.**	**Max**	**Min**
Reported Returns (in %)	0.917	4.935	17.059	-14.970
Unsmoothed Returns (in %)	0.915	5.606	18.976	-17.127

Panel B: Strategy Return Regression

	Dep. Var.: Avg. Hedge Fund Return
Constant	0.003**
	(2.641)
S&P 500	0.602**
	(14.793)
SC-LC	0.472**
	(8.749)
R2 (in %)	**76.19**

Panel C: Equity Long-Short Strategy Returns

	Average	**Std. Dev.**	**Max**	**Min**
Returns (in %)	0.826	3.481	11.194	-12.246

B and C of Table 3.1.

The motivation behind regressing the average returns on the two factors which proxy for the equity-long short strategy, according to Fung and Hsieh (2004), is that the selected funds are not exclusively investing in their core strategy. They might diversify their portfolios by pursuing non-core strategies. Regressing the returns on the S&P 500 and the SC-LC factors allows us to measure the return component which corresponds to the equity long-short strategy. Further, the exposure to the core strategy is probably correlated across funds: there are periods when the great majority of the hedge funds increase or decrease their risk exposure. The fitted returns based on the S&P 500 and the SC-LC factors proxy for the returns of an average equity long-short hedge fund, which holds its exposure to the equity long-short strategy constant over the total sample period. However, as a robustness check in Section 3.5, I use the equally weighted average returns across all 1,205 hedg funds as the equity long-short strategy returns, and my findings do not change.

3.3 Methodology

I distinguish between two channels through which an equity long-short hedge fund manager can change the risk of her portfolio. One way is to increase the exposure to the equity long-short strategy. The other way is to take on more risk through the non-core strategies. In the existing literature on risk-shifting by mutual funds and hedge funds (see, for example, Brown et al. (1996) and Aragon and Nanda (2012)), the variance of portfolio returns is used as a proxy for portfolio risk. My goal is to decompose the portfolio variance, so that we can distinguish between the two channels.

The exposure of hedge fund i's return in period t, $r_{i,t}$, to the return of the equity long-short strategy s in period t, $r_{s,t}$, is given by the parameter $\beta_{i,t}$ in the model:

$$r_{i,t} = \alpha_i + \beta_{i,t} r_{s,t} + \epsilon_{i,t},\ \epsilon_{i,t} \sim N(0, \tilde{\sigma}^2_{i,t}). \tag{3.3.1}$$

When $\beta_{i,t}$ increases, the hedge fund takes more exposure to the strategy, i.e. the hedge fund increases the equity long-short risk component of the portfolio (see Patton and Ramadorai (2013) for evidence of time-varying hedge fund risk exposure).

The $\tilde{\sigma}^2_{i,t}$ is the component of the return variance not due to the core strategy. This can be seen by decomposing the portfolio return variance as

$$Var(r_{i,t}) = \beta^2_{i,t} Var(r_{s,t}) + \tilde{\sigma}^2_{i,t}, \tag{3.3.2}$$

or more compactly:

$$\sigma^2_{i,t} = \beta^2_{i,t} \sigma^2_{s,t} + \tilde{\sigma}^2_{i,t}, \tag{3.3.3}$$

where $Var(r_{i,t}) = \sigma^2_{i,t}$ and $Var(r_{s,t}) = \sigma^2_{s,t}$. Hence, increasing risk by taking on more exposure to a non-core strategy is measured by $\tilde{\sigma}^2_{i,t}$.[1]

The main difficulty when trying to estimate the model parameters and the variances for each month t is the frequency of the available data and the length of the time series. If hedge fund return data were available at a daily frequency, as mutual fund return data, a monthly estimator of $\beta_{i,t}$ could simply be based on the realised daily returns. However, the great majority of hedge fund databases provide return data at a monthly frequency. Further, the time series of returns is relatively short, as reliable data are unavailable before January 1994.

How can we estimate the model with monthly data? This section describes two methods which can be used.

3.3.1 Rolling window estimation

A simple approach is to estimate the model parameters and the variances with rolling overlapping windows. For example, if the rolling window is set at 12 months,

[1] If a hedge fund uses the non-core strategies to hedge the exposure to the equity long-short strategy, an omitted variable bias results when estimating $\beta_{i,t}$. However, all of the hedge funds selected load strongly on the equity long-short strategy over the whole sample period. This implies that hedging through non-core strategies is very limited.

then for each month t, the model in equation 3.3.1 is estimated with observations from $t-5$ to $t+6$. However, this method, albeit simple to implement, suffers from several disadvantages.

First, the window length is ad-hoc, and the window length matters. Very different estimates of $\beta_{i,t}$ and $\tilde{\sigma}^2_{i,t}$, $\hat{\beta}_{i,t}$ and $\hat{\tilde{\sigma}}^2_{i,t}$, can result by changing the length of the rolling window. Generally, a short window will lead to a volatile and noisy estimate, while a long window results in an estimate which changes sluggishly. Thus, when investigating the drivers of $\hat{\beta}_{i,t}$ or $\hat{\tilde{\sigma}}^2_{i,t}$, the conclusions will be affected by the choice of window length. Since the *right* window length is unknown, an empirical analysis will fail to be conclusive.

Second, any empirical analysis with $\hat{\beta}_{i,t}$ or $\hat{\tilde{\sigma}}^2_{i,t}$ is further complicated due to the measurement error in the estimates. When including $\hat{\beta}_{i,t}$ as a regressor, an instrument is needed to deal with the endogeneity problem caused by the measurement error. Due to the overlapping windows, the measurement error will be serially correlated, which makes the use of lags, e.g. $\hat{\beta}_{i,t-1}$, as instruments problematic.

Third, for long windows, the estimates can be very persistent. For example, in regressions where $\hat{\beta}_{i,t}$ is the dependent variable and $\hat{\beta}_{i,t-1}$ is included as a regressor, unit root problems are likely to arise.

Fourth, a rolling window cannot estimate $\beta_{i,t}$ and $\tilde{\sigma}^2_{i,t}$ for every month in the sample – data points will be lost at the beginning and end of the sample. Hence, an already short time series becomes even shorter.

3.3.2 Dynamic conditional estimation

An alternative approach, which overcomes the short comings of the rolling window estimator for the purpose of measuring the time-varying hedge fund portfolio risk, is rooted in the conditional variance and correlation estimation literature. The conditional variances, $\sigma^2_{i,t}$ and $\sigma^2_{s,t}$, are estimated with models of the ARCH family (see Engle (1982)). As an estimator for $\beta_{i,t}$, the dynamic conditional β introduced

by Engle (2012) is used. The dynamic conditional β is denoted $\beta_{i,t}^{DC}$ for hedge fund i and month t and is given by

$$\beta_{i,t}^{DC} = \frac{\sigma_{i,t}\sigma_{s,t}\rho_{is,t}}{\sigma_{s,t}^2} = \frac{\sigma_{is,t}}{\sigma_{s,t}^2}. \tag{3.3.4}$$

The $\sigma_{i,t}$ and the $\sigma_{s,t}$ are the conditional volatilities of $r_{i,t}$ and $r_{s,t}$, respectively, which can be estimated using GARCH models. The $\rho_{is,t}$ is the conditional correlation of the two returns, which makes $\sigma_{is,t}$ the conditional covariance. The $\rho_{is,t}$ is estimated with the dynamic conditional correlation model, DCC model, introduced by Engle and Sheppard (2001) and Engle (2002). With the $\hat{\beta}_{i,t}^{DC}$ and the estimates of the conditional volatilities at hand, $\hat{\tilde{\sigma}}_{i,t}^{DC}$ is given by rearranging equation (3.3.3).

Compared to the rolling windows, this method enjoys the considerable advantage that the whole sample is utilised, which reduces noise, while time-variation is still permitted. Further, no ad-hoc choice of window length is made: the procedure is entirely data-driven.

3.3.2.1 Estimation of the conditional variances

A standard GARCH(1,1) model as described in Bollerslev (1987) estimates $\sigma_{i,t}^2$ by maximising the likelihood of the model:

$$r_{i,t} = \mu_i + v_{i,t},\ \sigma_{i,t}^2 = \omega + a_i v_{i,t-1}^2 + b_i \sigma_{i,t-1}^2, \tag{3.3.5}$$

$$v_{i,t} = \sigma_{i,t} e_{i,t},\ \text{and}\ e_{i,t} \overset{i.i.d.}{\sim} N(0,1), \tag{3.3.6}$$

where $\omega > 0$, $a_i \geq 0$, and $b_i \geq 0$, to ensure that the conditional variance is positive. However, estimating the above model for each hedge fund is not feasible due to the relatively short time series of hedge fund returns. For the rare case, where a hedge fund has reported returns for every month from January 1994 to December 2012, a maximum of 228 time series data points are available. To estimate a GARCH(1,1) model, this number is insufficient. The solution to this problem comes in the form of

a panel GARCH(1,1) model. Pakel et al. (2011) show that a panel of returns with a short time series, small T, but a large cross-section, large N, can be used to estimate the conditional variance of each return series by restricting the dynamic parameters of a GARCH(1,1) model to be the same across the panel. When using their method for my unbalanced panel of hedge fund returns, the panel GARCH(1,1) is specified as

$$r_{i,t} = \mu_i + u_{i,t},\ \sigma^2_{i,t} = \gamma_i(1-a-b) + au^2_{i,t-1} + b\sigma^2_{i,t-1}, \tag{3.3.7}$$

$$u_{i,t} = \sigma_{i,t}e_{i,t},\ \text{and}\ e_{i,t} \overset{i.i.d.}{\sim} N(0,1), \tag{3.3.8}$$

where $\gamma_i > 0$, $a \geq 0$, $b \geq 0$, and $a + b < 1$. Note that the dynamic parameters of the volatility process, a and b, are the same for all N hedge funds in the panel. The intercept term is allowed to vary across hedge funds due to the nuisance parameter γ_i.

The model parameters γ_i, a, and b, are estimated via a two stage procedure. Due to the GARCH(1,1) specification in equation (3.3.7), which is known as variance-tracking, γ_i is given by

$$\bar{\sigma}^2_i = \gamma_i, \tag{3.3.9}$$

where $\bar{\sigma}^2_i$ is the unconditional variance. This equality can be derived by taking unconditional expectations on both sides of the GARCH(1,1) process in (3.3.7). Based on this equality, γ_i can be estimated in the first step of the estimation procedure using a method of moments estimator:

$$E[u^2_{i,t}] = \gamma_i. \tag{3.3.10}$$

In the second stage of the estimation procedure, a and b are estimated by maximising

the composite likelihood for the $N \times T$ panel of hedge fund returns r:

$$CL(\psi_{(N_T)}; r) = \frac{1}{T} \sum_{t=1}^{T} \left[\frac{1}{N_t} \sum_{i=1}^{N_t} logf(r_{i,t} | \mathcal{F}_{t-1}; \psi_i) \right], \tag{3.3.11}$$

where $\psi_{N_T} = (\theta', \gamma'_{N_T}, \mu'_{N_T})'$, $\theta = (a, b)'$, $\gamma_{N_T} = (\gamma_1, \gamma_2, ..., \gamma_N)'$, $\mu_{N_T} = (\mu_1, \mu_2, ..., \mu_N)'$, $\mathcal{F}_{t-1}$ is the information set at time $t-1$, and N_t is the number of hedge funds with return data in period t (N_t differs across periods as the panel is unbalanced). The estimates of γ_{N_T} are taken from the first stage and are not estimated in the second stage.[2] In Section 3.5, I show that my results are robust to different specifications of the volatility model.

The conditional volatility of the equity long-short strategy returns, $\sigma_{s,t}$, can be estimated with a standard GARCH(1,1) model. As described in Section 3.2.2, the strategy return series is computed as the fitted returns of the model, which regresses the equally weighted average hedge fund returns on a constant and the two asset-based style factors S&P 500 and the SC-LC. Since data for the two factors are available for a long time series, the strategy return series can be extended. This ensures that enough observations are available to estimate $\sigma_{s,t}$ with the model:

$$r_{s,t} = \mu_s + u_{s,t},\ \sigma^2_{s,t} = \phi + a_s u^2_{s,t-1} + b_s \sigma^2_{s,t-1}, \tag{3.3.12}$$

$$u_{s,t} = \sigma_{s,t} e_{s,t},\ \text{and}\ e_{s,t} \overset{i.i.d.}{\sim} N(0,1), \tag{3.3.13}$$

via a quasi-maximum likelihood estimation with the restrictions $\phi > 0$, $a_s \geq 0$, and $b_s \geq 0$.

3.3.2.2 Estimation of the conditional correlation

The conditional correlation, $\rho_{is,t}$, between the returns of hedge fund i and the equity long-short strategy s is estimated with the DCC model described in Engle and

[2] This two stage estimation is subject to the "incidental parameter" problem when T is small, i.e. around a 100 time series observations, as shown by Pakel et al. (2011).

Sheppard (2001) and Engle (2002). A conditional correlation matrix for the two return series is modelled as

$$R_{is,t} = Q^*_{is,t} Q_{is,t} Q^*_{is,t}, \tag{3.3.14}$$

$$\text{where } Q_{is,t} = \begin{pmatrix} q_{ii,t} & q_{si,t} \\ q_{is,t} & q_{ss,t} \end{pmatrix} \text{ and } Q^*_{is,t} = \begin{pmatrix} 1/\sqrt{q_{ii,t}} & 0 \\ 0 & 1/\sqrt{q_{ss,t}} \end{pmatrix}. \tag{3.3.15}$$

The dynamics of $Q_{is,t}$ are given by

$$Q_{is,t} = (1 - c - d)\bar{R}_{is} + c(w_{is,t} w'_{is,t}) + dQ_{is,t-1}, \tag{3.3.16}$$

where $\bar{R}_{is}$ is the correlation matrix of the hedge fund return and the strategy return over the total sample. The 2×1 vector $w_{is,t} = (u_{i,t}/\sigma_{i,t},\, u_{s,t}/\sigma_{s,t})'$ contains the standardised residuals. As in the panel GARCH(1,1) above, the dynamic parameters of the model are restricted to be the same across all the hedge funds in the panel. A motivation for this restriction is given in Engle, Shephard, and Sheppard (2008).

3.3.2.3 Empirical tests to detect risk-shifting

To analyse, whether the performance of a hedge fund relative to a benchmark has an effect on risk taking in subsequent months, I use the estimates of $\beta_{i,t}$ and $\tilde{\sigma}_{i,t}$ as proxies for the risk exposure to the equity long-short strategy and non-core strategies, respectively. If there is evidence of risk-shifting in the data, we expect that hedge funds change $\beta_{i,t}$ or $\tilde{\sigma}_{i,t}$ based on their past performance relative to their peers. For example, a fund which is underperforming increases its exposure to the equity long-short strategy, $\beta_{i,t}$, to catch up. Arguably, a spike in $\tilde{\sigma}_{i,t}$ could also be due to an increase in the variance of the non-core strategies of fund i. In this case, the fund manager is not actively seeking greater exposure to the non-core strategies. However, this is noise in the estimates which makes it more difficult to detect risk-shifting. I test for risk-shifting among equity long-short hedge funds with a fund

fixed effects panel model specified as

$$Y_{i,t} = \pi_i + \lambda^+ I^+_{i,t-L}(r^{CR}_{i,t-L} - \bar{r}^{CR}_{t-L}) + \lambda^- I^-_{i,t-L}(r^{CR}_{i,t-L} - \bar{r}^{CR}_{t-L}) + \Gamma_1 Cont._{i,t-L} + \Gamma_2 Y_{i,t-L} + \eta_{i,t}, \tag{3.3.17}$$

where $Y_{i,t}$ is either $\beta_{i,t}$ or $\tilde{\sigma}_{i,t}$. The lag of the regressors in months is denoted L. Normal error terms, $\eta_{i,t}$, correlated across hedge funds in each time period are assumed. The variable which measures the performance of a fund relative to its peers is $(r^{CR}_{i,t-L} - \bar{r}^{CR}_{t-L})$. In line with the existing literature, I assume that hedge funds compare their year-to-date performance to their peers. Thus, $r^{CR}_{i,t-L}$ is the cumulative annual return for hedge fund i in month $t-L$, and $\bar{r}^{CR}_{t-L}$ is the average cumulative annual return across all of the equity long-short hedge funds in the sample, i.e. the return of the peer index. The dummy variable $I^+_{i,t}$ is equal to one if hedge fund i has been outperforming its peers during the calendar year up to month $t-L$ and zero otherwise. The opposite holds for the dummy variable $I^-_{i,t}$. The lag of the dependent variable is included as a regressor, as the risk exposure is expected to be relatively persistent due to liquidity, market impact, and transaction costs. As Aragon and Nanda (2012), I include hedge fund flows and the return autocorrelation as controls. Both are estimated over a 12 month rolling window. The flows are included to control for changes in portfolio composition. New funds which have to be redeployed or assets which have to be sold due to investor redemption can both change the risk of the portfolio without any active risk-shifting occurring. The autocorrelation of hedge fund returns controls for the possibility that the illiquidity of assets in the hedge fund portfolio drives the risk. A high autocorrelation is associated with an illiquid portfolio (see, for example, Getmansky et al. (2004) and Chapter 2 of this thesis).

3.4 Results

This section presents the results of the methodology outlined in Section 3.3. First, I show the results for the case where $\beta_{i,t}$ and $\tilde{\sigma}_{i,t}$ are estimated with rolling windows and highlight the issues which arise when using this approach. Then I present the results for the case where conditional volatilities and correlations are used instead.

3.4.1 Results for the rolling window estimation

Table 3.2 shows the results for the fund fixed effects panel model in equation (3.3.17) for $\beta_{i,t}$ and $\tilde{\sigma}_{i,t}$ estimated with rolling windows, as outlined in Section 3.3.1. The standard errors of the coefficients are adjusted for contemporaneous correlation using a delete-cross-section jackknife method (see Shao and Wu (1989) and Shao (1989) for examples of the implementation of the jackknife estimator). The estimate of λ^-, $\hat{\lambda}^-$, is inconsistent in terms of sign across varying window lengths. The coefficients are insignificant except when $\tilde{\sigma}_{i,t}$ is the dependent variable and the window length is 6 or 18 months. The estimate of λ^+, $\hat{\lambda}^+$, is at most weakly significant when positive but strongly significant when negative. The reason behind the inconsistent results can be seen in Figures 3.1 and 3.2. The estimates of $\beta_{i,t}$ and $\tilde{\sigma}_{i,t}$ change substantially across different window lengths. Further, the estimates of the 6 month window are far more volatile (standard deviation of 0.33) than those of the 12 month window (standard deviation of 0.22).

Since the *right* window length is unknown, these results are of little benefit in providing conclusive insights into the relation between time-varying risk exposure and past performance. An additional concern regarding the results is that, for longer windows, the dependent variables are highly persistent and close to a unit root. A Dickey-Fuller test on the time series of the equally weighted average of $\tilde{\sigma}_{i,t}$ estimated over a 12 and 18 month window fails to reject the unit root hypothesis (including a time trend and constant, combined and separately, does not change the outcome

of the test, the smallest p-value is 0.15 for the 12 month window and 0.19 for the 18 month window). Hence, the regression results in the last two columns of Table 3.2 are spurious. When differencing the model to address the unit root problem, $\hat{\lambda}^{+}$ and $\hat{\lambda}^{-}$ fail to be significant.

3.4.2 Results for the dynamic conditional estimation

Alternatively, the estimators of the dynamic conditional estimation method, $\hat{\beta}_{i,t}^{DC}$ and $\hat{\tilde{\sigma}}_{i,t}^{DC}$, can be used to measure the exposure to the core and non-core strategies. The parameter estimates of the panel GARCH(1,1) model in equation (3.3.7) are reported in Panel A of Table 3.3. The estimates for the volatility process of the equity long-short strategy in equation (3.3.12) are shown in Panel B. The parameter estimates of the DCC model are given in Panel C. For the standard GARCH(1,1) model, the Bollerslev-Wooldridge covariance estimator is applied to estimate the standard errors (see Bollerslev and Wooldridge (1992)). For the Panel GARCH(1,1) model, Pakel et al. (2011) show how the covariance of the parameter estimates can be computed by taking the estimation of the nuisance parameter in the first stage into account. The standard errors of the DCC model are based on a block bootstrap with 500 iterations – I sample the funds and use moving time series blocks with 114 (half of the total sample) observations of the standardised residuals. All of the parameter estimates are significant, which supports the volatility model specifications. In Section 3.5, I change the the specification of the volatility models and find that $\hat{\beta}_{i,t}^{DC}$ and $\hat{\tilde{\sigma}}_{i,t}^{DC}$ are robust.

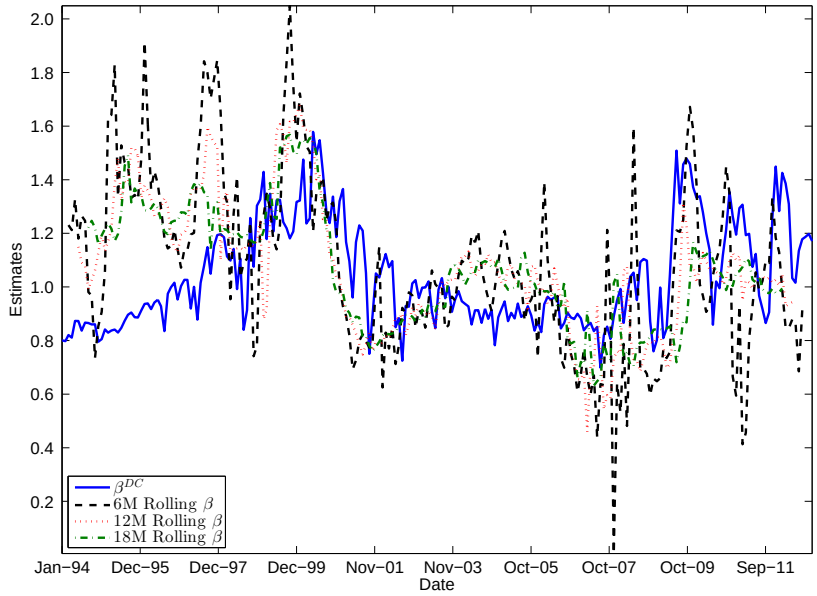
Figure 3.1 plots the estimated time-varying exposure to the equity-long short strategy, $\hat{\beta}_{i,t}^{DC}$, averaged across all hedge funds. During the Dot-com bubble and in the aftermath of the financial crisis of 2008, the exposure increased.

Table 3.2: Panel Model With Rolling Window Estimates

Reported are the results of the fund fixed effects panel regression. The estimates of $\beta_{i,t}$ and $\tilde{\sigma}_{i,t}$ act as dependent variables. These estimates are based on rolling windows. Results are shown for various window lengths. The data are monthly from January 1994 to December 2012. The jackknife standard errors are clustered by time. Significant coefficients at the 10% and 5% level are indicated by * and **, respectively.

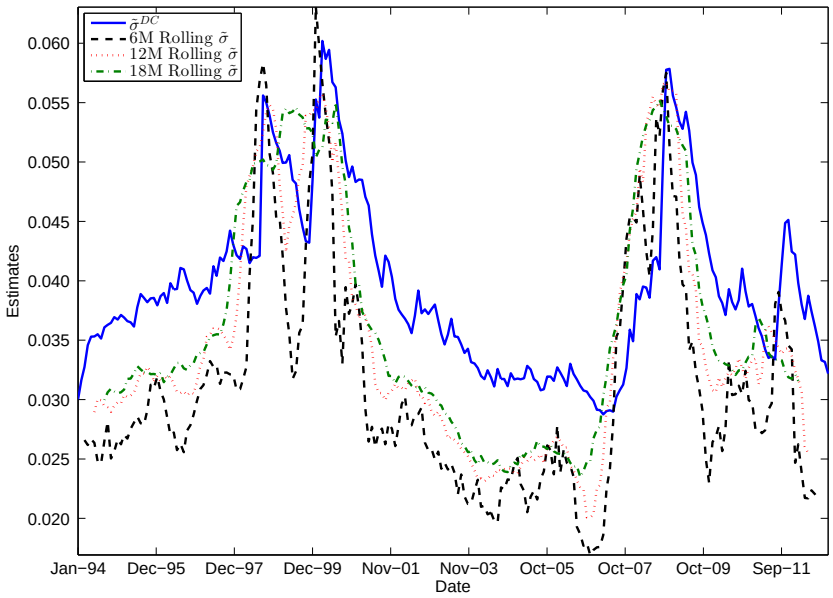
	Dep. Variable: $\hat{\beta}_{it}$			**Dep. Variable:** $\hat{\tilde{\sigma}}_{i,t}$		
Window Length	6M	12M	18M	6M	12M	18M
$I^{+}_{i,t-1} \times (r^{CR}_{i,t-1} - \bar{r}^{CR}_{t-1})$	0.191*	0.070*	0.060*	0.712*	0.010	-0.158**
	(1.860)	(1.863)	(1.703)	(1.772)	(0.130)	(-3.428)
$I^{-}_{i,t-1} \times (r^{CR}_{i,t-1} - \bar{r}^{CR}_{t-1})$	-0.044	-0.003	0.045	-0.973**	-0.092	0.159*
	(-0.275)	(-0.050)	(1.380)	(-3.443)	(-0.671)	(1.691)
$\text{Flows}_{\text{i,t-1}}$	0.026	0.009	0.000	-0.052	-0.029	-0.008
	(1.108)	(0.717)	(0.020)	(-1.067)	(-0.981)	(-0.379)
$\text{Autocorr.}_{\text{i,t-1}}$	0.079**	0.013	0.005	0.029	-0.023	-0.003
	(2.861)	(0.970)	(0.671)	(0.618)	(-0.971)	(-0.191)
$\text{Dep. Var.}_{\text{i,t-1}}$	0.779**	0.919**	0.950**	0.855**	0.950**	0.974**
	(30.728)	(79.006)	(98.908)	(74.152)	(138.205)	(169.625)
R2 (in %)	**61.06**	**85.22**	**90.71**	**74.15**	**90.23**	**94.27**
Observations	**90,440**	**87,619**	**84.790**	**90,440**	**87,619**	**84.790**

Figure 3.1: Average β_t Estimates over Time



This figure plots the average β_t estimate across all 1,205 equity long-short hedge funds for the total sample period, January 1994 to December 2012.

Figure 3.2: Average $\tilde{\sigma}_t$ Estimates over Time



This figure plots the average $\tilde{\sigma}_t$ estimate across all 1,205 equity long-short hedge funds for the total sample period, January 1994 to December 2012.

An increase in β_t during times when the stock market rallies is reasonable, as most equity long-short funds have a positive exposure to stocks – as evident from the S&P 500 return used as the asset-based style factor by Fung and Hsieh (2004) to explain the returns of equity long-short funds. Further, we can observe that $\hat{\beta}_{i,t}^{DC}$ is correlated with the rolling estimates: the correlation with the 6, 12, and 18 month estimates, is 0.27, 0.30, and 0.28, respectively. However, the correlation is relatively weak. Hence, there is a substantial discrepancy between the series. Below, we see that this discrepancy matters for the estimation of the panel model in (3.3.17). Figure 3.2 shows the portfolio risk due to non-core strategies averaged across all of the hedge funds over time. The measure spikes during the crash of the Dot-com bubble and the financial crisis. The rolling estimates correlate with $\hat{\tilde{\sigma}}_{i,t}^{DC}$ more strongly than in the case of the $\hat{\beta}_{i,t}^{DC}$: the correlations are 0.73, 0.79, and 0.80, for a window length of 6, 12, and 18 months, respectively, but the point estimates are quite different for some periods.

The results for the panel regression in (3.3.17) with $\hat{\beta}_{i,t}^{DC}$ or $\hat{\tilde{\sigma}}_{i,t}^{DC}$ as dependent variables are provided in Table 3.4. $\hat{\lambda}^{+}$ and $\hat{\lambda}^{-}$ show consistent signs across all lags. I use lags of one to three months, as this lag length takes into account that changing the risk of a portfolio is a slow process due to liquidity, market impact, and transactions costs. The benchmarking variables are strongly significant in every case. The negative significant $\hat{\lambda}^{-}$ implies that hedge funds which underperform their peers year-to-date, tend to increase their risk in subsequent months through both the equity long-short strategy and non-core strategies. This finding is in line with Brown et al. (2001) and Aragon and Nanda (2012), who find that underperforming hedge funds engage in risk-shifting. Interestingly, it is not only the underperfomers that increase their risk: hedge funds that outperform their peers also tend to take on more risk in subsequent periods, implied by the strongly significant and positive $\hat{\lambda}^{+}$. An potential explanation for this phenomenon is suggested by the behavioural finance literature. For example, Johnson and Thaler (1990) and Barberis, Huang,

Table 3.3: Volatility Estimates

Panel A reports the volatility parameter estimates for the hedge fund return panel estimated with a panel GARCH(1,1). The estimates and t-stats for the nuisance parameter γ_i are averaged across the panel. Newey-West standard errors are used for the nuisance parameter. The hedge fund return data starts in January 1994 and ends in December 2012. Panel B reports the volatility parameter estimates for the equity long-short strategy returns estimated with a GARCH(1,1) model. The time series for the equity long-short strategy starts in January 1979 and ends in December 2012. The parameter estimates of the DCC model are given in Panel C. The standard errors are estimated with a block bootstrap. The DCC model is estimated with data from January 1994 to December 2012. Significant coefficients at the 10% and 5% level are indicated by * and **, respectively.

Panel A: Panel GARCH(1,1) Volatility Parameter Estimates for Hedge Fund Returns

	Average γ_i	a	b
Parameter Estimates	0.004**	0.124**	0.805**
T-Stats	4.578	4.651	23.017

Panel B: GARCH(1,1) Volatility Parameter Estimates for Strategy Returns

	ϕ	a_s	b_s
Parameter Estimates	0.000**	0.217**	0.316*
T-Stats	3.277	3.194	1.873

Panel C: Parameter Estimates for DCC Model

	c	d
Parameter Estimates	0.025**	0.931**
T-Stats	2.304	35.421

and Santos (2001), point out that investors with prior gains become less risk averse and are prone to make riskier investments in the future.

Note that the estimate of λ^- is greater than the estimate of λ^+ in all six cases. This implies that when simply analysing which funds increase risk the most, the result will be: the losers. However, I find that not only the losers increase their risk. The winners do so, too, but by less.

Since $\hat{\beta}_{i,t}^{DC}$ and $\hat{\tilde{\sigma}}_{i,t}^{DC}$ are both quite persistent, I have tested for a unit root using the Dickey-Fuller test on the time series of the equally weighted average of $\hat{\beta}_{i,t}^{DC}$ and $\hat{\tilde{\sigma}}_{i,t}^{DC}$ across all hedge funds. The test strongly rejects the hypothesis of a unit root – I include a constant and a time trend. The same holds for the augmented Dickey-Fuller test.

A defining feature in the mutual fund tournament literature is the convex relationship between performance and fund flows (see, for example, Chevalier and Ellison (1997) and Sirri and Tufano (1998)). The better a mutual fund performs relative to its peers, the more fund inflows it will receive in the future – the fund flows are thought to be the main motivation behind the risk-shifting of mutual funds. In the hedge fund literature, the link between flows and performance is less well documented, as differences in lock-up periods and hedge funds being closed to new investors complicate the analysis. Nonetheless, a positive correlation between fund flows and performance can be expected, which implies that one factor behind the risk-shifting of hedge fund managers are investor flows. Thus, an interesting question is if hedge funds which are less sensitive to fund outflows due to longer redemption periods manage their risk differently than hedge funds with a short redemption period. We can expect that a hedge fund with a long redemption period is less dependent on retaining the trust of investors with returns above the industry average in the short term to avoid fund outflows. The opposite holds for a hedge fund with a short redemption period. Hence, empirical tests should reveal more pronounced risk-shifting behaviour for hedge funds with short redemption periods.

To test this empirically, I include the benchmarking variables interacted with a redemption dummy variable in the panel regression. The dummy variable takes the value one if the hedge fund has a long redemption period. A long redemption period is defined as being above the 75 percentile of all 1,205 funds in the sample. In Table 3.5, the results are presented. Hedge funds with long redemption periods engage less in risk-shifting when outperforming and underperforming. The coefficients of the benchmarking variables interacted with the redemption dummy variable are strongly significant and of the opposite sign compared to the coefficient of the benchmarking variables that are not interacted with the redemption dummy. This holds for the exposure to the equity long-short strategy and to the non-core strategies. We can infer that hedge funds with longer redemption periods are less pressured by potential redemptions and this weakens the link between past performance and subsequent risk taking.

One might expect that the HWM is also a relevant benchmark for hedge funds. The HWM determines whether a performance fee can be charged – if an HWM exists, a hedge fund's AUM has to be above the HWM in order to charge a performance fee. Hence, a hedge fund manager with assets below the HWM has a strong incentive to increase returns to push the AUM above the HWM. However, Brown et al. (2001) find that the performance relative to the HWM is not relevant for risk-shifting. Aragon and Nanda (2012) conclude that hedge funds with an HWM are more conservative and engage less in risk shifting than hedge funds without an HWM. I follow Aragon and Nanda (2012) and track a hedge funds HWM for each year y, $HWM_{i,y}$, with its annual returns, $r_{i,y}$:

$$AUM_{i,y} = AUM_{i,y-1} \times (1 + r_{i,y}), \tag{3.4.1}$$

$$HWM_{i,y} = max[HWM_{i,y-1}, AUM_{i,y}], \tag{3.4.2}$$

where $AUM_{i,0} = 1$ and $HWM_{i,0} = 1$. The benchmarking variable in the panel

model is then defined as $(r^{CR}_{i,t-L} - HWM_{i,y})$: the difference between the year-to-date return and the HWM of the corresponding year. The panel regression estimates are reported in Table 3.6. For hedge funds below the HWM, there is very little evidence of risk-shifting. Hedge funds which are above the HWM seem to engage in some risk-shifting, but the significance of the coefficients is far weaker than when comparing the year-to-date return to the peer index. Also the magnitude of the coefficients is substantially lower. Hence, my estimates of time-varying risk support the finding that the HWM is less relevant than the peer index for the benchmarking of hedge funds.

The evidence for the risk-shifting of outperforming and underperforming funds leads to the question of whether the risk taking behaviour differs across the year. To answer this question, I add a quarterly dummy variable interacted with the benchmarking variable to the panel model. The dummy variable takes the value one if the month of the lagged benchmarking variable, $t-1$, belongs to a particular quarter. Only one quarterly dummy variable at a time is added to the model. The results are presented in Table 3.7. For outperforming funds, the dummy variables which correspond to the 1$^{\text{st}}$ and the 3$^{\text{rd}}$ quarter, respectively, have a significant coefficient for risk exposure to the core and non-core strategies. For the 1$^{\text{st}}$ quarter dummy, the sign of the coefficient is positive. This implies that the risk-shifting based on the performance in the 1$^{\text{st}}$ quarter is stronger than that during the rest of the year. However, for the 3$^{\text{rd}}$ quarter dummy variable, the sign flips. Hedge funds which are ahead of their peers appear to lock-in their gains and reduce the risk in their portfolios. Unfortunately, the results for underperforming funds are less insightful. The coefficients on the quarterly dummy variables fail to be significant in all but one case.

Table 3.4: Panel Model With Dynamic Conditional Estimates

Reported are the results of a fund fixed effects panel regression with estimates of $\beta_{i,t}$ and $\tilde{\sigma}_{i,t}$ as dependent variables. The $\hat{\beta}_{i,t}^{DC}$ and the corresponding $\hat{\tilde{\sigma}}_{i,t}^{DC}$ are used as estimates. Results are shown with regressors lagged by one, two, and three months. The data are monthly from January 1994 to December 2012. The jackknife standard errors are clustered by time. Significant coefficients at the 10% and 5% level are indicated by * and **, respectively.

	Dep. Variable: $\hat{\beta}_{i,t}^{DC}$			**Dep. Variable:** $\hat{\tilde{\sigma}}_{i,t}^{DC}$		
Lags of Regressors	L=1	L=2	L=3	L=1	L=2	L=3
$I_{i,t-L}^{+} \times (r_{i,t-L}^{CR} - \bar{r}_{t-L}^{CR})$	0.180**	0.284**	0.250**	0.962**	0.959**	1.140**
	(4.189)	(4.882)	(3.040)	(4.708)	(3.641)	(3.289)
$I_{i,t-L}^{-} \times (r_{i,t-L}^{CR} - \bar{r}_{t-L}^{CR})$	-0.193**	-0.379**	-0.423**	-0.977**	-1.184**	-1.485**
	(-2.514)	(-4.242)	(-4.072)	(-4.798)	(-5.120)	(-5.712)
$\text{Flows}_{\text{i,t-L}}$	-0.000	-0.004	-0.006	-0.004	-0.015	-0.026
	(-0.002)	(-0.368)	(-0.455)	(-0.153)	(-0.481)	(-0.716)
$\text{Autocorr.}_{\text{i,t-L}}$	0.012	0.027	0.037**	0.072**	0.060	0.041
	(0.964)	(1.443)	(2.040)	(2.303)	(1.511)	(0.817)
$\text{Dep. Var.}_{\text{i,t-L}}$	0.851**	0.744**	0.647**	0.894**	0.811**	0.725**
	(35.131)	(25.810)	(20.699)	(84.009)	(53.137)	(41.029)
R2 (in %)	**74.87**	**59.79**	**45.95**	**84.93**	**70.96**	**58.41**
Observations	**93,205**	**92,298**	**91,372**	**93,205**	**92,298**	**91,372**

Table 3.5: Panel Model With Dynamic Conditional Estimates and Redemption Dummy

Reported are the results of a fund fixed effects panel regression with estimates of $\beta_{i,t}$ and $\tilde{\sigma}_{i,t}$ as dependent variables. The $\hat{\beta}_{i,t}^{DC}$ and the corresponding $\hat{\tilde{\sigma}}_{i,t}^{DC}$ are used as estimates. The benchmarking variables are interacted with a dummy variable, I_i^{LR}, which takes the value one when the length of the redemption period of the hedge fund is above the 75th percentile across all funds and zero otherwise. The data are monthly from January 1994 to December 2012. The jackknife standard errors are clustered by time. Significant coefficients at the 10% and 5% level are indicated by * and **, respectively.

	Dependent Variable:	
	$\hat{\beta}_{i,t}^{DC}$	$\hat{\tilde{\sigma}}_{i,t}^{DC}$
$I_{i,t-1}^{+} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	0.194**	1.049**
	(4.301)	(4.810)
$I_{i,t-1}^{-} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	-0.208**	-1.037**
	(-2.681)	(-4.651)
$I_i^{LR} \times I_{i,t-1}^{+} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	-0.078**	-0.456**
	(-2.896)	(-5.103)
$I_i^{LR} \times I_{i,t-1}^{-} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	0.100**	0.297**
	(2.084)	(2.096)
$\text{Flows}_{\text{i,t-1}}$	-0.000	-0.007
	(-0.044)	(-0.279)
$\text{Autocorr.}_{\text{i,t-1}}$	0.012	0.072**
	(0.948)	(2.345)
$\text{Dep. Var.}_{\text{i,t-1}}$	0.852**	0.895**
	(35.266)	(82.703)
R2 (in %)	**75.22**	**85.09**
Observations	**87,424**	**87,424**

Table 3.6: Panel Model With Dynamic Conditional Estimates and HWM Benchmark

Reported are the results of a fund fixed effects panel regression with estimates of $\beta_{i,t}$ and $\tilde{\sigma}_{i,t}$ as dependent variables. The $\hat{\beta}_{i,t}^{DC}$ and the corresponding $\hat{\tilde{\sigma}}_{i,t}^{DC}$ are used as estimates. Results are shown with regressors lagged by one, two, and three months. The HWM of a fund for the corresponding year is used as a performance benchmark. The data are monthly from January 1994 to December 2012. The jackknife standard errors are clustered by time. Significant coefficients at the 10% and 5% level are indicated by * and **, respectively.

	Dep. Variable: $\hat{\beta}_{i,t}^{DC}$			**Dep. Variable:** $\hat{\tilde{\sigma}}_{i,t}^{DC}$		
Lags of Regressors	L=1	L=2	L=3	L=1	L=2	L=3
$I_{i,t-L}^{+} \times (r_{i,t-L}^{CR} - HWM_{i,y})$	0.023*	0.045**	0.047**	0.150*	0.156*	0.199*
	(1.792)	(2.091)	(1.971)	(1.871)	(1.717)	(1.683)
$I_{i,t-L}^{-} \times (r_{i,t-L}^{CR} - HWM_{i,y})$	-0.004	-0.008	-0.016	-0.039**	-0.029	-0.016
	(-0.487)	(-0.755)	(-1.333)	(-2.540)	(-1.262)	(-0.585)
$\text{Flows}_{\text{i,t-L}}$	-0.003	-0.013	-0.016	0.014	-0.002	-0.008
	(-0.252)	(-1.065)	(-1.211)	(0.581)	(-0.056)	(-0.229)
$\text{Autocorr.}_{\text{i,t-L}}$	0.020	0.044**	0.061**	0.093**	0.097**	0.092
	(1.437)	(2.149)	(3.057)	(2.630)	(2.106)	(1.566)
$\text{Dep. Var.}_{\text{i,t-L}}$	0.872**	0.775**	0.681**	0.922**	0.848**	0.722**
	(38.447)	(27.634)	(21.850)	(93.539)	(63.933)	(51.039)
R2 (in %)	**76.94**	**61.59**	**48.06**	**85.89**	**72.62**	**60.27**
Observations	**54,583**	**54,081**	**53,570**	**54,583**	**54,081**	**53,570**

3.5 Robustness Checks

This section shows that the results in the previous section are robust for different specifications of the panel model in (3.3.17), modifications to the volatility models in (3.3.7) and (3.3.12), differencing the variables, and variations of the computation of the equity long-short strategy returns.

The estimates $\hat{\beta}_{i,t}^{DC}$ and $\hat{\tilde{\sigma}}_{i,t}^{DC}$ are probably measured with error. Hence, when including the lag of the dependent variable as a regressor in the panel model, the measurement error could lead to an endogeneity problem. As the dependent variables are persistent, lags can be used as instruments, under the assumption that the measurement error is not serially correlated. Table 3.8 shows the results when using the second lag of the dependent variable as an instrument for the first lag. The benchmarking variables remain strongly significant, and the coefficient signs are consistent with the results reported in Table 3.4. Adding additional lags as instruments does not alter the results.

I also use the year-to-date return of the hedge fund as a benchmarking variable and ignore the distance to the year-to-date return of the peer index. The coefficient signs do not change. The t-statistics are slightly lower but remain significant, except in one case, where the t-statistic is not significant.

As a third robustness check, the sensitivity to the specification of the volatility models is analysed. I replace the GARCH(1,1) models in (3.3.7) and (3.3.12) with AVGARCH(1,1) models (see Taylor (1986)). In an AVGARCH(1,1) model, the conditional volatility and not the conditional variance is directly modelled. The exact specification is shown in Appendix 3.A. The results are insensitive to this change. The magnitude and significance of the benchmarking variables' coefficients are very similar to the GARCH(1,1) case.

I also estimate the panel model in (3.3.17) after differencing all the variables. The signs of the λ^+ and λ^- estimates remain the same. The estimates are also strongly significant except for λ^- when $\hat{\tilde{\sigma}}_{i,t}^{DC}$ is the dependent variable.

Table 3.7: Panel Model With Dynamic Conditional Estimates and Quarterly Dummy

Reported are the results of a fund fixed effects panel regression with estimates of $\beta_{i,t}$ and $\tilde{\sigma}_{i,t}$ as dependent variables. The $\hat{\beta}_{i,t}^{DC}$ and the corresponding $\hat{\tilde{\sigma}}_{i,t}^{DC}$ are used as estimates. The benchmarking variables are interacted with a dummy variable, I_t^Q, which takes the value one, if the month of the regressor lies within a particular quarter. The quarter corresponding to each regression is given at the top of each column. The data are monthly from January 1994 to December 2012. The jackknife standard errors are clustered by time. Significant coefficients at the 10% and 5% level are indicated by * and **, respectively.

Panel A: Dependent Variable $\hat{\beta}_{i,t}^{DC}$

Quarter Dummy	1st Quarter	2nd Quarter	3rd Quarter	4th Quarter
$I_{i,t-1}^{+} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	0.169**	0.170**	0.240**	0.162**
	(4.241)	(3.677)	(4.585)	(2.141)
$I_{i,t-1}^{-} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	-0.220**	-0.163*	-0.280**	-0.135
	(-2.902)	(-1.877)	(-3.346)	(-1.029)
$I_{t-1}^{Q} \times I_{i,t-1}^{+} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	0.512*	0.186	-0.269**	0.020
	(1.795)	(1.263)	(-2.554)	(0.218)
$I_{t-1}^{Q} \times I_{i,t-1}^{-} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	-0.004	-0.304	0.306	-0.103
	(-0.013)	(-1.137)	(1.406)	(-0.597)
$\text{Flows}_{i,t-1}$	0.001	-0.000	-0.000	-0.000
	(0.079)	(-0.020)	(-0.008)	(-0.001)
$\text{Autocorr.}_{i,t-1}$	0.011	0.012	0.014	0.013
	(0.893)	(0.933)	(1.096)	(1.010)
$\hat{\beta}_{i,t-1}^{DC}$	0.847**	0.848**	0.851**	0.852**
	(35.315)	(36.141)	(35.221)	(35.053)
R2 (in %)	**75.02**	**75.01**	**75.17**	**74.89**
Observations	**93,205**	**93,205**	**93,205**	**93,205**

Panel B: Dependent Variable $\hat{\tilde{\sigma}}_{i,t}^{DC}$

Quarter Dummy	1st Quarter	2nd Quarter	3rd Quarter	4th Quarter
$I_{i,t-1}^{+} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	0.941**	0.977**	1.104**	1.137**
	(4.435)	(4.476)	(4.855)	(3.648)
$I_{i,t-1}^{-} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	-1.166**	-0.970*	-0.901**	-1.296**
	(-5.750)	(-4.334)	(-4.726)	(-3.792)
$I_{t-1}^{Q} \times I_{i,t-1}^{+} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	3.787**	-0.194	-0.652**	-0.221
	(2.501)	(-0.581)	(-2.319)	(-0.515)
$I_{t-1}^{Q} \times I_{i,t-1}^{-} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	-1.516**	0.035	-0.208	0.531
	(-2.980)	(0.087)	(-0.392)	(1.364)
$\text{Flows}_{i,t-1}$	-0.001	-0.004	-0.004	-0.004
	(-0.023)	(-0.144)	(-0.159)	(-0.169)
$\text{Autocorr.}_{i,t-1}$	0.067**	0.073**	0.072**	0.069**
	(2.048)	(2.301)	(2.348)	(2.182)
$\hat{\tilde{\sigma}}_{i,t-1}^{DC}$	0.881**	0.895**	0.895**	0.892**
	(97.256)	(84.352)	(86.415)	(88.199)
R2 (in %)	**85.56**	**84.93**	**85.00**	**84.97**
Observations	**93,205**	**93,205**	**93,205**	**93,205**

The last robustness check is with regard to the computation of the equity long-short strategy returns. The strategy returns for the baseline case are the fitted values when regressing the equally weighted average returns of all 1,205 hedge funds on the asset-based style factors S&P 500 and SC-LC. In columns 6 and 11 of Table 3.8, we can see that the results are robust to simply using the equally weighted average returns instead of the fitted returns.

3.6 Conclusion

Measuring the risk of hedge fund portfolios is important for regulators and investors. Since information on hedge fund holdings is highly restricted, researchers use the variance of returns as a proxy for risk. However, most hedge funds report returns at a monthly frequency and reliable data are unavailable before January 1994. Hence, estimating a time-varying variance is a challenge due to the short time series and low frequency. In this chapter, I propose a new method for measuring the time-varying variance of hedge fund returns for each month. I focus on equity long-short hedge funds, but my methodology can be adapted to other investment strategies. The method is based on recent advances in the conditional variance and correlation estimation literature. In addition to measuring the time-varying risk, it allows me to decompose the variance into three components: the exposure to the core-strategy, i.e. the equity long-short strategy, the variance of the core strategy, and the variance of the non-core strategies. Compared to estimating and decomposing the return variance with OLS estimations over rolling overlapping windows, the approach in this chapter enjoys considerable advantages. The main advantage is that an ad-hoc choice of window lengths – estimates can be substantially different across window lengths – is avoided.

After estimating and decomposing the time-varying risk, I investigate whether the risk taking of equity long-short hedge funds is driven by the past performance

Table 3.8: Robustness Checks

Reported are the results of a panel regression with estimates of $\beta_{i,t}$ and $\tilde{\sigma}_{i,t}$ as dependent variables, respectively. The $\hat{\beta}_{i,t}^{DC}$ and the corresponding $\hat{\tilde{\sigma}}_{i,t}^{DC}$ are used as estimates. Results are shown for different benchmarking variables, variations in the lag of the dependent variable, two volatility models (the GARCH(1,1) and the AVGARCH(1,1)), for differenced variables, and for variations in the computation of the strategy returns. "Fitted" indicates that the mean returns across all hedge funds are regressed on the two factors S&P 500 and SC-LC, and the fitted returns are used as the strategy returns. "Average" indicates that the mean returns across all hedge funds are used as the strategy returns. Fund fixed effects are used for the variables in levels. The data are monthly from January 1994 to December 2012. The jackknife standard errors are clustered by time. Significant coefficients at the 10% and 5% level are indicated by * and **, respectively.

	Dep. Variable: $\hat{\beta}_{i,t}^{DC}$					**Dep. Variable:** $\hat{\tilde{\sigma}}_{i,t}^{DC}$				
Volatility Model	GARCH	GARCH	AVGARCH	GARCH	GARCH	GARCH	GARCH	AVGARCH	GARCH	GARCH
Variables	Levels	Levels	Levels	Differences	Levels	Levels	Levels	Levels	Differences	Levels
Strategy Returns	Fitted	Fitted	Fitted	Fitted	Avg.	Fitted	Fitted	Fitted	Fitted	Avg.
$I_{i,t-1}^{+} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	0.289**		0.148**	0.286**	0.056**	1.530**		0.763**	1.826**	0.822**
	(4.940)		(3.648)	(2.556)	(1.998)	(4.181)		(4.666)	(4.054)	(5.971)
$I_{i,t-1}^{-} \times (r_{i,t-1}^{CR} - \bar{r}_{t-1}^{CR})$	-0.319**		-0.168**	-0.145**	-0.064*	-1.777**		-0.815**	-0.763	-0.866**
	(-3.621)		(2.301)	(-2.281)	(-1.697)	(-6.564)		(-5.029)	(-1.054)	(-6.144)
$I_{i,t-1}^{+} \times (r_{i,t-1}^{CR})$		0.152**					0.778**			
		(3.658)					(3.494)			
$I_{i,t-1}^{-} \times (r_{i,t-1}^{CR})$		-0.207					-1.738**			
		(-1.607)					(-4.437)			
$\text{Flows}_{\text{i,t-1}}$	0.002	-0.001	-0.002	0.027	-0.003	-0.020	0.009	-0.005	0.027	0.004
	(0.195)	(-0.113)	(-0.153)	(1.084)	(-0.571)	(-0.662)	(0.318)	(-0.256)	(0.524)	(0.170)
$\text{Autocorr.}_{\text{i,t-1}}$	0.021	0.010	0.013	-0.007	0.004	0.055	0.050*	0.064**	0.212**	0.060**
	(1.225)	(0.802)	(1.081)	(-0.199)	(1.079)	(1.044)	(1.858)	(2.397)	(2.111)	(2.740)
$\text{Dep. Var.}_{\text{i,t-1}}$		0.849**	0.834**	-0.151**	0.908**		0.881**	0.911**	0.000	0.893**
		(33.412)	(36.606)	(-2.852)	(139.373)		(65.351)	(88.798)	(0.025)	(110.126)
$\text{Dep. Var.}_{\text{i,t-2}}$	0.751**					0.803**				
	(25.677)					(53.919)				
R2 (in %)	**59.95**	**74.95**	**71.34**	**3.80**	**83.10**	**72.17**	**85.40**	**86.80**	**7.74**	**84.80**
Observations	**93,205**	**93,205**	**93,205**	**91,991**	**93,205**	**93,205**	**93,205**	**93,205**	**91,991**	**93,205**

of a hedge fund relative to the performance of its peers. I find that the exposure to the core strategy and the variance due to non-core strategies are higher for hedge funds which have strongly underperformed or outperformed in recent months. The risk-shifting for underperforming funds is stronger in terms of magnitude. That funds which try to catch up with their peers increase the risk of their portfolios is a well established result in the mutual fund and hedge fund tournament literature. However, the fact that outperforming hedge funds take on more exposure to risky strategies is a new finding. This phenomenon could potentially be explained by prior gains leading to less risk aversion, as found in the behavioural finance literature.

I also discover strong evidence that the risk-shifting is mitigated for hedge funds with long redemption periods. This could imply that hedge funds with long redemption periods are less concerned with fund outflows due to poor performance in the short term, and thus, the need to achieve high returns through an increase in risk is attenuated. Further, the HWM does not appear to be a benchmark based on which risk-shifting occurs. My results are robust across variations in the specification of panel regression models, volatility processes, and computation of the equity long-short strategy returns.

I view the time-varying risk measure methodology developed in this chapter as a useful addition to the hedge fund literature. In future work, I aim to analyse the implications of losers and winners increasing risk simultaneously for the hedge fund industry as a whole but also for the asset markets in general.

Appendix

3.A The Panel AVGARCH(1,1) Model

The AVGARCH(1,1) model for the panel of hedge fund returns is specified as

$$r_{i,t} = \mu_i + \tau_{i,t},\ \sigma_{i,t} = \gamma_{i,AVG}(1 - a_{AVG}\kappa_i - b_{AVG}) + a_{AVG}|\tau_{i,t-1}| + b_{AVG}\sigma_{i,t-1}, \tag{3.A.1}$$

$$\tau_{i,t} = \sigma_{i,t}e_{i,t},\ \text{and}\ e_{i,t} \overset{i.i.d.}{\sim} N(0,1), \tag{3.A.2}$$

where the κ_i is given by $E[|\epsilon_{i,t}|]/\bar{\sigma}_i = \kappa_i$, with $\bar{\sigma}_i$ being the unconditional standard deviation. The restrictions needed to ensure that $\sigma_{i,t}$ is positive are: $\gamma_{i,AVG} > 0$, $a_{AVG} \geq 0$, $b_{AVG} \geq 0$, and $a_{AVG}\kappa_i + b_{AVG} < 1$. In the first stage of the estimation procedure, the estimators for $\gamma_{i,AVG}$ and κ_i are $\sqrt{E[\epsilon_{i,t}^2]} = \gamma_{i,AVG}$ and $E[|\epsilon_{i,t}|]/\sqrt{E[\epsilon_{i,t}^2]} = \kappa_i$, respectively. The second stage estimation is the same as for the panel GARCH(1,1) model outlined in Section 3.3.2.1.

For the equity long-short strategy returns, the AVGARCH(1,1) model is

$$r_{s,t} = \mu_s + \tau_{s,t},\ \sigma_{s,t} = \phi_{AVG} + a_{s,AVG}|\tau_{s,t-1}| + b_{s,AVG}\sigma_{s,t-1}, \tag{3.A.3}$$

$$\tau_{s,t} = \sigma_{s,t}e_{s,t},\ \text{and}\, e_{s,t} \overset{i.i.d.}{\sim} N(0,1). \tag{3.A.4}$$

Chapter 4

Model-Based Priors for Predicting the Equity Premium

Abstract

The out-of-sample predictability of the equity premium is strongly debated. This chapter attempts to improve the forecast performance of typical single variable predictive regressions used in the equity premium prediction literature through Bayesian priors derived from consumption-based asset pricing models. To implement the model-based priors, I develop a Bayesian procedure which is rooted in the macroeconometrics literature. The model-based priors are derived from four asset pricing models: Habit Formation, Habit Formation Term Structure, Long Run Risk, and Prospect Theory. I find that the model-based priors can increase the explanatory power, measured by the out-of-sample R^2, of the single variable predictive regressions by several percentage points.

4.1 Introduction

Predicting aggregate stock returns has been of great interest to finance practitioners and academic finance economists alike. There exists an extensive literature that uses a variety of variables to predict stock market returns. Several papers, such as Campbell and Shiller (1988) and Fama and French (1988), found that valuation ratios are correlated with subsequent returns. Other variables used to predict the equity premium are yields on corporate bonds and treasuries (see, for example, Campbell (1987) and Fama and French (1989)). More recently, new predictor variables have been suggested, for example, by Baker and Wurgler (2000) and Lettau and Ludvigson (2001). These papers use data on corporate payout and financing activity and the level of consumption relative to wealth, respectively.

Goyal and Welch (2008), henceforth Goyal and Welch, provide a comprehensive analysis of the in-sample and out-of-sample (OOS) predictive power of all the major variables used in the literature up to that point. They estimate a single variable predictive regression in-sample and out-of-sample (OOS) – the equity premium is predicted OOS by estimating the parameters with data up to t and forecasting the equity premium in $t+1$:

$$Equity\ Premium_{t+1} = \beta_0 + \beta_1 x_t + \epsilon_{t+1},\ \text{where } \epsilon_{t+1} \sim N(0, \sigma^2), \tag{4.1.1}$$

where the equity premium, defined as the rate of return on the stock market minus the prevailing short term interest rate, is regressed on a lagged predictor x_t. Goyal and Welch estimate this model for various predictors, for different time periods, and at multiple frequencies. Their findings are rather sobering: while most of the predictors are not even significant in-sample, those that have some predictive power in-sample still suffer from very poor OOS performance. Goyal and Welch compare the predictive regressions with the historical average, i.e. predicting that the equity premium in $t+1$ will be what it has been on average up to t, i.e. $\beta_1 x_t$ is eliminated

from the model. The historical average outperforms the predictive regression OOS for the majority of the predictors. Hence, an investor would have been better off using the historical average to forecast the equity premium rather than any of the available variables.

Campbell and Thompson (2008), henceforth Campbell and Thompson, respond to the claims of Goyal and Welch by imposing restrictions on the parameter estimates of the predictive regressions. Through these restrictions, the equity premium can be forecast OOS such that a real time investor could profit. Pettenuzzo, Timmermann, and Valkanov (2013) also show how economic constraints on forecasting models result in considerable economic gains when forecasting the equity premium. This chapter attempts to further investigate the effects of economic constraints on model parameters. I address the claim of poor OOS predictability of the equity premium by implementing restrictions through well specified Bayesian priors. Consumption-based asset pricing models act as a source for these priors – hence, I call them "model-based priors" – which are then imposed on the single variable predictive regression. The procedure can be considered as estimating the parameters in equation (4.1.1) with a data set that combines empirical data with data simulated from the asset pricing model. I find that the model-based priors can raise the OOS R^2 by several percentage points for all of the predictor variables considered. The results are particularly good for the postwar data sample starting in 1947.

Besides trying to overcome the weak predictability of the equity premium OOS, my approach can also be interpreted as a different perspective on the assessment of asset pricing models. Generally, consumption-based asset pricing models are evaluated based on their ability to match empirical data moments: in this chapter, the effectiveness of their priors is scrutinised.

The asset pricing models used to derive the model-based priors share a common origin. They have been developed to offer a solution to the equity premium puzzle (see Mehra and Prescott (1985)). I use four models to derive model-based priors.

These models are the Habit Formation (HF) model (see Campbell and Cochrane (1999)), the Prospect Theory (PT) model (see Barberis, Huang, and Santos (2001)), the Long Run Risk (LRR) model (see Bansal and Yaron (2004)), and the Habit Formation Term Structure (HFTS) model (see Wachter (2006)). I briefly describe these models in Appendix 4.B.

The structure of this chapter is as follows. In Section 4.2, I explain the motivation behind the Bayesian method. Section 4.3 discusses the Bayesian procedure. In Section 4.4, I describe the empirical and simulated data used for my analysis. Section 4.5 presents the results. Section 4.6 discusses the working of the model-based prior in detail, and Section 4.7 concludes.

4.2 Bayesian Motivation

The Bayesian method is related to Del Negro and Schorfheide (2004), who combine a dynamic stochastic general equilibrium (DSGE) model with a vector autoregression (VAR) and create a DSGE-VAR. I use prior distributions derived from consumption-based asset pricing models to estimate the parameters of the single variable predictive regression in equation (4.1.1). These model-based priors can improve the OOS forecasts of the equity premium in two ways. First, they reduce the overfitting of the coefficient estimates. The signals of the predictors are generally very weak, and thus, parameters estimated solely on empirical data suffer from the influence of noise. The idea of Bayesian priors improving the forecast performance of a regression model by reducing the overfitting dates back to the Minnesota prior of Doan, Litterman, and Sims (1984). Second, because of the iterative nature of the OOS forecasting procedure, the OOS estimates of β_0 and β_1, denoted $\hat{\beta}_0$ and $\hat{\beta}_1$, respectively, can be very different from their in-sample counterparts at some points in the sample. In these cases, the fundamental relationship between the predictor and the equity premium, which holds over the total sample, fails to hold temporarily.

This leads inevitably to inaccurate predictions of the equity premium. Depending on the theoretical estimates implied by the asset pricing model, the model-based prior can bring the parameter estimates closer to their fundamental values during such periods – the model-based prior helps to bring the OOS performance closer to the in-sample performance without introducing an in-sample look-ahead bias.

The data sets used by the authors to calibrate the asset pricing models overlap partially with my OOS forecasting period. Hence, there exists an implicit look-ahead bias. To address this issue, I also use the model-based priors to predict the equity premium over a subsample of the data which has no overlap with the calibration data, and the results are robust to this adjustment of the sample period.

How does my methodology compare to that of Campbell and Thompson? Campbell and Thompson impose three restrictions on the parameter estimates of the predictive regression sequentially – where restrictions 2 and 3 only apply to valuation ratios. First, whenever $\hat{\beta}_1$ takes on a sign which is not in line with the in-sample estimate, it is set to zero. Second, they set the intercept to zero whenever it is negative and the coefficient to one when greater than one and zero when negative. Third, they set $\hat{\beta}_0$ equal to zero and $\hat{\beta}_1$ equal to one regardless of the estimates. For each of the three restrictions, Campbell and Thompson successfully improve the OOS forecasts of the single variable predictive regressions. The third restriction is particularly interesting: Campbell and Thompson set the constant and the coefficient to zero and one, respectively, because these are the theoretical values implied by the steady state theory of Gordon (1962). This can be considered as imposing a model-based prior very dogmatically, i.e. the prior is very tight and no empirical data are used to estimate β_0 and β_1.[1]

This chapter builds on Campbell and Thompson and develops their parameter

[1]Campbell and Thompson also improve the OOS forecast performance through changing the predictor variable by adding a growth term derived from the Gordon (1962) growth model or restricting the equity premium forecast to be greater or equal to zero. Since I solely focus on the parameter estimation of the predictive regression rather than on transforming the predictor or imposing sign restrictions on the forecast, I do not discuss this aspect of their paper.

restrictions further. My methodology differs from theirs as the model-based priors are based on a well specified Bayesian procedure. This has the advantage of allowing us to adjust the tightness of the prior gradually rather than imposing it dogmatically. Further, the model-based prior can reduce the overfitting of parameter estimates which do not violate any of the boundaries, i.e. an estimate of β_0 which lies above zero or a coefficient estimate which is between zero and one. The model-based prior can also push negative estimates above zero and positive estimates below one, where the restrictions used by Campbell and Thompson only set them as equal to zero or one, respectively. The two approaches are compared in Table 4.8 in Appendix 4.C.

The model-based prior is rooted in the economic theory of the consumption based asset pricing models. An alternative prior is the Minnesota prior, introduced by Doan et al. (1984), which is a purely statistical way of reducing overfitting in parameter estimates. A comparison of the performance of the model-based priors with an adaption of the Minnesota prior is performed in Section 4.5.7. This comparison makes it possible to assess whether the economic theory of the asset pricing models helps to improve the OOS forecasting performance beyond the statistical shrinking of the parameter estimates by the Minnesota prior.

4.3 Econometric Methodology

This section lays out the econometric methodology. First, I describe the empirical procedure for forecasting the equity premium and the assessment of the forecast performance. Second, I illustrate the Bayesian procedure.

4.3.1 Empirical procedure

The predictive regressions are specified as in Goyal and Welch and Campbell and Thompson. The equity premium is denoted by y_t and is defined as the rate of return on the stock market minus the prevailing short term interest rate. The

equity premium is regressed on a constant and an independent variable lagged by one period, the predictor x_{t-1}:

$$y_t = \beta_0 + \beta_1 x_{t-1} + \epsilon_t, \text{ where } \epsilon_t \overset{i.i.d.}{\sim} N(0, \sigma^2). \tag{4.3.1}$$

For a sample which starts in period 0 and ends in period T, the in-sample regression uses all available data, i.e. T data points, to run the regression in (4.3.1). The OOS predictions of the equity premium are generated through recursive forecasts. I take all available observations up to period t and estimate the model in (4.3.1). Based on the resulting estimate $\hat{\beta}$ of the 2×1 vector $\beta = [\beta_0, \beta_1]'$ and the observation x_t, one can forecast the equity premium in $t+1$. The predicted equity premium is denoted by $\hat{y}^p_{t+1}$. Although I have observations up to T, observations after t are not used to estimate β: hence, a real time investor who is interested in predicting the equity premium can implement this procedure. Since at the beginning of the sample there is an insufficient number of observations to estimate β accurately, the recursive OOS forecast starts after period z where $z < T$, i.e. observations from period 1 to z are used as burn-in data.[2]

To measure the OOS forecast performance, I compare the predictive regression forecasts with predictions based on the historical average model. The historical average model forecasts the equity premium in $t+1$, denoted by $\hat{y}^h_{t+1}$, as its mean up to period t – we conjecture that the best forecast for the equity premium in $t+1$ is what it has been on average up to t. This is equivalent to estimating the regression in (4.3.1) with β_1 set equal to zero.

The two models are compared via an OOS R^2, which is computed as

$$R^2_{OS} = 1 - \frac{\sum_{t=z+1}^{T}(y_t - \hat{y}^p_t)^2}{\sum_{t=z+1}^{T}(y_t - \hat{y}^h_t)^2}. \tag{4.3.2}$$

[2]I also attempted a rolling forecast procedure, which does not use the whole sample to estimate β. Instead, a number of periods j is determined and the regression in (4.3.1) is estimated based on observations from $t-j$ to t. The model-based prior is also successful for rolling windows. However, the recursive forecast performs generally better than the rolling forecast.

In the denominator, not the average equity premium over the whole sample is used – as it would be for an in-sample R^2 – but the historical average equity premium $\hat{y}_t^h$. A high OOS R^2 is a sign that the predictive regression is the superior model. Campbell and Thompson emphasize the importance of the OOS R^2 statistic for a real time investor. They show that even a small positive OOS R^2 can lead to considerable utility gains and an increase in expected returns for investors with standard mean-variance preferences.

4.3.2 Model-based prior

This section describes the model-based prior. A detailed description of the Bayesian procedure and the derivation of the posterior estimates is provided in Appendix 4.A. The implementation of the prior is based on Del Negro and Schorfheide (2004).

The posterior distribution of the predictive regression parameters is

$$p(\beta, \sigma^2 \mid Y, X) = \frac{p(Y, X \mid \beta, \sigma^2) p(\beta, \sigma^2 \mid \theta)}{p(Y, X \mid \theta)}, \tag{4.3.3}$$

where Y is a $T \times 1$ vector with elements y_t, X is a $T \times 2$ matrix with rows $[1\, x_t]$, β and σ^2 are the parameters of the single variable predictive regression, and θ stands for the parameters of the consumption-based asset pricing model. The parameters of the asset pricing model are set equal to the values calibrated by the authors of the respective model. The prior for β and σ^2 derived from the asset pricing model is given by $p(\beta, \sigma^2|\theta)$. I specify the prior in two different ways. The first specification is called the population moments (PM) prior. When imposing the PM prior, the modes of the posterior distributions of β and σ^2 are given by

$$\tilde{\beta}_{PM}(\theta) = (\lambda T \Gamma_{xx}^*(\theta) + X'X)^{-1}(\lambda T \Gamma_{xy}^*(\theta) + X'Y) \text{ and} \tag{4.3.4}$$

$$\begin{aligned}\tilde{\sigma}^2_{PM}(\theta) =& \frac{1}{(\lambda+1)T}[(\lambda T\Gamma^*_{yy}(\theta) + Y'Y) \\ & - (\lambda T\Gamma^*_{yx}(\theta) + Y'X)(\lambda T\Gamma^*_{xx}(\theta) + X'X)^{-1}(\lambda T\Gamma^*_{xy}(\theta) + X'Y)],\end{aligned} \tag{4.3.5}$$

where $\lambda T\Gamma^*_{yy}(\theta)$, $\lambda T\Gamma^*_{yx}(\theta)$, and $\lambda T\Gamma^*_{xx}(\theta)$, with $\Gamma^*_{yy}(\theta) = E[y_t^2]$, are the population moments of the data simulated from the asset pricing model scaled by T and the hyperparameter λ, which determines the weight of the model-based prior – how strongly we believe the model implied parameter estimates relative to the empirical estimates. The population moments are estimated through a Monte Carlo procedure, discussed in Section 4.4.2.1. The estimators in (4.3.4) and (4.3.5), respectively, can be interpreted as OLS estimators for a data set that combines empirical and simulated data. The second specification of the prior is the empirical state variables (ESV) prior. The modes of the posterior distributions of β and σ^2 when imposing the ESV prior are

$$\tilde{\beta}_{ESV}(\theta) = (\lambda\bar{X}'\bar{X}(\theta) + X'X)^{-1}(\lambda\bar{X}'\bar{Y}(\theta) + X'Y) \text{ and} \tag{4.3.6}$$

$$\begin{aligned}\tilde{\sigma}^2_{ESV}(\theta) =& \frac{1}{(\lambda+1)T}[(\lambda\bar{Y}'\bar{Y}(\theta) + Y'Y) \\ & - (\lambda\bar{Y}'\bar{X}(\theta) + Y'X)(\lambda\bar{X}'\bar{X}(\theta) + X'X)^{-1}(\lambda\bar{X}'\bar{Y}(\theta) + X'Y)].\end{aligned} \tag{4.3.7}$$

The state variables of all the asset pricing models considered in this chapter depend on at least one of the three variables: consumption growth, dividend growth, and inflation rate. The asset pricing models have assumed data generating processes for these variables, which are used to simulate the model randomly. For example, the HF model specifies the consumption growth process as $\Delta c_{t+1} = g + v_{t+1}$, where g is a constant and $v_{t+1} \overset{i.i.d.}{\sim} N(0, \sigma_v^2)$. However, instead of drawing random shocks v_{t+1} and simulating a consumption growth series, one can feed the empirical consumption growth data into the model. The resulting time series of the predictor and the equity premium are then used for the model-based ESV prior. The model data simulated with empirical state variables are denoted by $(\bar{Y}, \bar{X})$. As with the population

moments, the simulated data can be scaled by λ to give them more weight in the posterior distribution.

The PM and the ESV prior can be interpreted as two different ways of assessing the asset pricing models. The PM prior requires that the theoretical relationship between the equity premium and the predictive variable implied by the model provides useful information for estimating the coefficients. Hence, the model is assessed regarding the empirical validity of its theory. The ESV prior is a more difficult challenge for the asset pricing model. The ESV prior also assesses the theoretical relationship between the equity premium and the predictor implied by the simulated data, but the simulated data are based on empirical consumption growth, dividend growth, or inflation, and not on the model's assumed data generating processes.

The procedure laid out above is based on the total sample of size T. For a recursive forecasting exercise, I must implement the model-based prior in every period t with data up to t. The resulting posterior coefficient estimates, given in equations (4.3.4) and (4.3.6), are then used to forecast the equity premium in $t+1$.

For the model-based prior to be of any value for a real time investor, the weight of the prior has to be estimated OOS. Hence, λ is estimated by maximising the data density $p(Y, X \mid \theta)$ in equation (4.3.3) numerically over a grid with data up to period z and held constant for the whole forecasting exercise.[3] Generally speaking, the λ value which maximises the data density is higher when the simulated data match the empirical data more closely. The exact forms of the data densities for the PM and ESV prior are given in Appendix 4.A.3.

4.4 Data

The availability of variables which can be used to implement the model-based priors is determined by the four asset pricing models. The variables of the equity premium prediction literature for which at least one model provides theoretical predictions are:

[3]Reestimating the λ before every forecast changes the results only marginally.

the equity premium, the dividend-price ratio, the dividend yield, the long term yield, the short term yield, and the term spread. The equity premium is the difference between the log stock returns and the log of the risk-free rate. The dividend-price ratio is computed as the difference between the log dividends and the log price. The dividend yield is defined as the difference between the log dividends and the log price lagged by one period. The long term yield is the annualised rate on long term U.S. government bonds. The short term yield is given by the annualised rate on the short term U.S. government debt. The term spread is the difference between the two.

4.4.1 Empirical data

The empirical time series of the equity premium and the predictor variables are taken from Amit Goyal's website. This ensures that I use the same data as Goyal and Welch did, with the only difference being that my sample ends in 2011 instead of 2005. The returns and price of the S&P 500 index are provided by Amit Goyal from 1872 to 2011. The original source of the time series up to 1926 is Robert Shiller's website. From 1926, returns from The Center for Research in Security Prices are used. The risk-free rate is the 3-month U.S. Treasury bill rate. Dividends on the S&P 500 index are 12-month moving sums, and the time series starts in 1871 and ends in 2011. For the yields, the time series of the annualised rates on the 3-month U.S. Treasury bill and the long term U.S. government bonds start in 1920 and 1919, respectively.

4.4.2 Simulated data

An overview of how I simulate data from each asset pricing model is given below. I simulate the asset pricing models at the frequency proposed by the respective authors.

4.4.2.1 Monte Carlo simulation

Three out of the four asset pricing models lack analytical solutions. Hence, I estimate the model-based population moments, which are needed to implement the PM prior, through a Monte Carlo simulation.[4] The set-up of the Monte Carlo simulation is the same for all models. I randomly simulate data (Y^*, X^*), where Y^* is a vector containing the simulated equity premium and X^* is a matrix containing a constant and the simulated predictor variable, based on the assumed data generating processes of the respective model for S periods and estimate the population moments $\Gamma^*_{yy}(\theta)$, $\Gamma^*_{yx}(\theta)$, and $\Gamma^*_{xx}(\theta)$, through the sample moments. The estimators are $\hat{\Gamma}^*_{yy}(\theta) = E[y_t^{*2}]$, $\hat{\Gamma}^*_{yx}(\theta) = E[y_t^*(1,\ x_{t-1}^*)]$, and $\hat{\Gamma}^*_{xx}(\theta) = E[(1, x_{t-1}^*)'(1, x_{t-1}^*)]$. This procedure is iterated P times, which results in a total of $S \times P$ observations, and the estimates of the population moments are averaged across the P simulations. Since the equity premium and the predictive variables generated by the models are covariance stationary, averaging the P estimates should result in a precise estimate of the population moments conditional on S and P being large enough.[5]

4.4.2.2 Habit Formation model

Campbell and Cochrane (1999) simulate the HF model at a monthly frequency and aggregate the simulated data to match annual moments of the empirical data. I follow them and use priors from the HF model at an annual frequency. The model generates data for the equity premium, the dividend-price ratio, and the dividend yield.[6]

For the PM prior, I implement the Monte Carlo procedure by simulating data for $10,000$ years of data, estimating the population moments, and iterating the

[4] The LRR model has approximate analytical solutions, but these and the Monte Carlo estimates are identical.

[5] I use the number of simulated periods by the authors of the respective asset pricing model as a guideline for my simulation.

[6] The HF model has two specifications. I use their baseline specification which assumes perfect correlation between consumption and dividend growth.

procedure 10 times.

To simulate data for the ESV prior, we must feed empirical data for real log consumption growth. Real log consumption growth is defined as log growth in real per capita consumption expenditure on non-durables and services in the U.S. The annual data are taken from John Campbell's website and extended with data from Federal Reserve Economic Data (FRED). The consumption data series from John Campbell's website starts in 1890 at an annual frequency. To simulate the model at a monthly frequency when feeding real log consumption growth, the annual consumption growth rate is divided into monthly consumption growth rates.

4.4.2.3 Prospect Theory model

The PT model by Barberis et al. (2001) is simulated at an annual frequency in their paper. To be consistent with the model calibration of the authors, I use priors from the PT model at an annual frequency. The model generates data for the equity premium, the dividend-price ratio, and the dividend yield.

To estimate the population moments for the PM prior, I simulate $10,000$ years of data and average the estimates over 10 iterations.

To simulate data for the ESV prior, only one empirical time series is required. This is the time series for the error terms of the real dividend growth process. Since I have the real dividend growth rate and the parameter values provided by Barberis et al. (2001), I can back out the error terms.[7]

4.4.2.4 Long Run Risk model

Similar to Campbell and Cochrane (1999), Bansal and Yaron (2004) simulate the model at a monthly frequency and aggregate the simulated data to match the annual empirical moments. Hence, priors from the LRR model are used at an annual

[7]Barberis et al. (2001) run their model with different parameter values for b_0 and k. I use the specification which is most successful in matching the empirically observed moments of asset returns. This specification sets b_0 equal to 100 and k equal to 8. More details are provided in Appendix 4.B.

frequency in this chapter. The simulated variables are the equity premium, the dividend-price ratio, and the dividend yield.

For the PM prior, the Monte Carlo procedure is implemented by estimating the population moments from simulated data for $100,000$ months aggregated to annual levels and averaging them over 100 iterations.

Real log consumption growth and real log dividend growth act as input variables to generate data for the ESV prior. Further, the two state variables, the long run risk component and the time varying volatility, are required. I follow Bansal, Kiku, and Yaron (2007), who provide a simple procedure for extracting the two state variables from the data. The consumption data are the same as for the HF model.

4.4.2.5 Habit Formation Term Structure model

Wachter (2006) simulates the HFTS model at a quarterly frequency and assesses the model against quarterly empirical moments. Thus, I use priors from the HFTS model at a quarterly frequency. The simulated variables are: the equity premium, the long term yield, the short term yield, and the term spread.[8]

For the PM prior, the population moments are estimated from $40,000$ quarters of simulated data and averaged over 10 iterations.

To implement the ESV prior, two empirical time series are required. The first one is real log consumption growth. FRED provides quarterly data on consumption expenditure starting in 1947. The second one is inflation. I compute inflation based on CPI values from FRED.

[8] I could also simulate data on the dividend-price ratio and the dividend yield, but the theoretical relationship between these variables and the equity premium is the same as implied by the HF model.

4.5 Results

4.5.1 Habit Formation model-based priors

The HF model generates priors for the dividend-price ratio and the dividend yield. The results of imposing these priors on predictive regressions with annual forecast horizons are shown in Table 4.1. Panel A shows the results for the PM prior, and Panel B presents the results for the ESV prior. For both priors, I forecast the equity premium with annual returns and overlapping annual returns. For annual returns, the forecasting procedure is exactly as described in Section 4.3.1, with t being a calender year. For overlapping annual returns, the forecasts are also at annual horizons but iterated on a monthly basis, with t comprising 12 months of data. The motivation behind overlapping annual returns is to increase the number of predictions. This is particularly important when only subsamples of the data set are used, as in Section 4.5.6.

The empirical OOS statistics are computed by running the predictive regression solely on empirical data. I show the results when imposing the model-based prior for a value of λ which is estimated as explained in Section 4.3.2. To estimate λ, I only use data up to and not including the first period of the forecasting exercise. The OOS R^2 assesses how the predictive regression with or without model-based prior performs relatively to the historical average. In the ΔR^2_{OS} column, the difference between the OOS R^2 of the predictor with and without model-based prior is shown – a positive value indicates that the model-based prior leads to an improvement in the OOS forecast.

I analyse the performance of both priors over two time periods. The first period covers the total sample. For the PM prior, only empirical return data and dividend data are required. Hence, the time period starts in 1872. For the ESV prior, we also need empirical consumption data which is fed to the model. The sample period starts in 1890 for the dividend-price ratio and 1891 for the dividend yield, since

annual consumption data are unavailable before this. The second sample period used for the analysis is the postwar period, from 1947 to 2011. When choosing the starting period of the OOS forecast, I use at least 20 years of data for the initial regression. However, I do not start the OOS exercise before 1926 since high quality S&P 500 return data from CRSP are available from 1926 onwards – this start date has been motivated by Campbell and Thompson.

The PM prior is successful in improving the OOS performance of the predictive regression for both predictors and sample periods regardless of whether annual returns or overlapping annual returns are used. The OOS R^2 are increased up to several percentage points. The ESV prior fares far worse for the total sample than for the postwar sample. This can be explained by the great challenge the HF model faces when trying to replicate the market movements of the great depression and the two world wars. However, the performance of the ESV prior for the postwar sample is remarkable. Improving the OOS R^2 from a negative value up to 6.96% through model-based data which is simulated by feeding an empirical consumption growth process is impressive.

For the PM prior, the values for λ are higher than for the ESV prior. This means that we put more weight on the simulated data when estimating the constant and coefficient of the predictive regression when imposing the PM prior. The increase in OOS R^2 might appear small. However, Campbell and Thompson show that even a small positive OOS R^2 can translate into substantially higher expected returns for an investor with mean-variance preferences. They derive the difference in expected returns for an investor who forecasts the equity premium with a predictor and an investor who relies on the historical average return model as

$$\frac{1}{\gamma}\left(\frac{R^2_{OS}}{1-R^2_{OS}}\right)(1+S^2), \tag{4.5.1}$$

where γ is the risk aversion of the mean-variance preferences, and S is the Sharpe

ratio of the equity premium. The Sharpe ratio of the annual equity premium from 1967 to 2011 was 0.231. An investor who uses the dividend-price ratio as a predictor and imposes the ESV prior of the HF model over the postwar sample achieves and OOS R^2 of 6.48%. Hence, the difference in annual expected returns in percentage points compared to an investor who relies on the historical average model is 7.30 for a risk aversion of 1 and 2.43 for a risk aversion of 3.

4.5.2 Prospect Theory model-based priors

The PT model provides priors for the dividend-price ratio and the dividend yield at an annual frequency. The results for implementing these model-based priors are shown in Table 4.2.

The PM prior improves the OOS R^2 regardless of the predictor, the data sample, and whether we use annual returns or annual overlapping returns.[9] The PM prior fares extremely well for the postwar sample by pushing the OOS R^2 up to 5%. In most cases, the λ values are smaller than 1, implying that the empirical data are dominant in the posterior, except for the dividend yield in the postwar sample. There, the simulated data are weighted much stronger in the posterior than the empirical data. This indicates that the PT model matches the empirical data well. The ESV prior performs poorly for the total sample and fails to improve the OOS performance for both predictor variables. However, for the postwar sample, the ESV prior improves the OOS predictability substantially, with the OOS R^2 being of similar magnitude as when using the PM prior.

4.5.3 Long Run Risk model-based priors

Through the LRR model, we can obtain priors for the dividend-price ratio and the dividend yield. The priors are implemented at an annual frequency and shown in

[9]Since the PT model is simulated at an annual frequency, I only use the PM prior for the overlapping forecast procedure and not the ESV prior.

Table 4.1: Results for the Habit Formation Model-Based Priors

Reported are the results for the model-based priors from the HF model. The weight of the prior, λ, is estimated as explained in Section 4.3.2. The ΔR^2_{OS} stands for the difference between the R^2_{OS} of the predictor with and without model-based prior. A positive value indicates an improvement in the OOS forecasts through the prior. The data sample ends in December 2011. For the overlapping annual returns, the forecasting procedure is repeated in monthly steps, but the forecasts are at annual horizons.

Panel A: PM Prior

	Total Sample			**Postwar Sample**		
Annual Returns	Sample Begin: 1872 Forecast Begin: 1926			Sample Begin: 1947 Forecast Begin: 1967		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio						
Empirical OOS Statistics:		-0.39%			-0.87%	
PM OOS Statistics:	0.87	-0.38%	0.01%	1.11	2.64%	3.51%
Predictor: Div. Yield						
Empirical OOS Statistics:		-0.35%			-0.10%	
PM OOS Statistics:	1.58	0.67%	1.02%	2.66	0.86%	0.96%
Overlap. Ann. Returns	Sample Begin: 1M 1872 Forecast Begin: 12M 1926			Sample Begin: 12M 1947 Forecast Begin: 12M 1967		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio						
Empirical OOS Statistics:		-1.41%			-0.15%	
PM OOS Statistics:	0.96	-0.25%	1.16%	0.84	3.86%	4.01%
Predictor: Div. Yield						
Empirical OOS Statistics:		-0.85%			-0.03%	
PM OOS Statistics:	1.02	0.74%	1.59%	1.81	2.62%	2.65%

Panel B: ESV Prior

	Total Sample			**Postwar Sample**		
Annual Returns	Sample Begin: 1890 DP / 1891 DY Forecast Begin: 1926			Sample Begin: 1947 Forecast Begin: 1967		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio						
Empirical OOS Statistics:		-0.28%			-0.87%	
ESV OOS Statistics:	0.08	-84.18%	-83.90%	0.73	6.48%	7.35%
Predictor: Div. Yield						
Empirical OOS Statistics:		0.13%			-0.10%	
ESV OOS Statistics:	0.08	-66.37%	-66.50%	0.41	2.50%	2.60%
Overlap. Ann. Returns	Sample Begin: 12M 1890 DP / 12M 1891 DY Forecast Begin: 12M 1926			Sample Begin: 12M 1947 Forecast Begin: 12M 1967		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio						
Empirical OOS Statistics:		-1.94%			-0.15%	
ESV OOS Statistics:	0.09	-83.42%	-81.48%	0.73	6.96%	7.11%
Predictor: Div. Yield						
Empirical OOS Statistics:		-0.61%			-0.03%	
ESV OOS Statistics:	0.09	-69.74%	-69.13%	0.52	2.89%	2.92%

Table 4.2: Results for the Prospect Theory Model-Based Priors

Reported are the results for the model-based priors from the PT model. The weight of the prior, λ, is estimated as explained in Section 4.3.2. The ΔR^2_{OS} stands for the difference between the R^2_{OS} of the predictor with and without model-based prior. A positive value indicates an improvement in the OOS forecasts through the prior. The data sample ends in December 2011. For the overlapping annual returns, the forecasting procedure is repeated in monthly steps, but the forecasts are at annual horizons.

Panel A: PM Prior

	Total Sample			**Postwar Sample**		
Annual Returns	Sample Begin: 1872 Forecast Begin: 1926			Sample Begin: 1947 Forecast Begin: 1967		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio						
Empirical OOS Statistics:		-0.39%			-0.87%	
PM OOS Statistics:	0.22	-0.07%	0.32%	0.57	5.44%	6.31%
Predictor: Div. Yield						
Empirical OOS Statistics:		-0.35%			-0.10%	
PM OOS Statistics:	0.24	0.63%	0.98%	15.32	4.96%	5.06%
Overlap. Ann. Returns	Sample Begin: 1M 1872 Forecast Begin: 12M 1926			Sample Begin: 12M 1947 Forecast Begin: 12M 1967		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio						
Empirical OOS Statistics:		-1.41%			-0.15%	
PM OOS Statistics:	0.27	-0.15%	1.26%	0.39	5.79%	5.94%
Predictor: Div. Yield						
Empirical OOS Statistics:		-0.85%			-0.03%	
PM OOS Statistics:	0.33	0.42%	1.27%	9.20	5.79%	5.82%

Panel B: ESV Prior

	Total Sample			**Postwar Sample**		
Annual Returns	Sample Begin: 1890 DP / 1891 DY Forecast Begin: 1926			Sample Begin: 1947 Forecast Begin: 1967		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio						
Empirical OOS Statistics:		-0.28%			-0.87%	
ESV OOS Statistics:	0.56	-0.60%	-0.32%	0.56	5.65%	6.52%
Predictor: Div. Yield						
Empirical OOS Statistics:		0.55%			-0.10%	
ESV OOS Statistics:	0.50	-0.40%	-0.95%	0.39	4.05%	4.15%

Table 4.3: Results for the Long Run Risk Model-Based Priors

Reported are the results for the model-based priors from the LRR model. The weight of the prior, λ, is estimated as explained in Section 4.3.2. The ΔR^2_{OS} stands for the difference between the R^2_{OS} of the predictor with and without model-based prior. A positive value indicates an improvement in the OOS forecasts through the prior. The data sample ends in December 2011. For the overlapping annual returns, the forecasting procedure is repeated in monthly steps, but the forecasts are at annual horizons.

Panel A: PM Prior

	Total Sample			**Postwar Sample**		
Annual Returns	Sample Begin: 1872 Forecast Begin: 1926			Sample Begin: 1947 Forecast Begin: 1967		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio						
Empirical OOS Statistics:		-0.39%			-0.87%	
PM OOS Statistics:	5.05	1.47%	1.86%	1.90	5.52%	6.39%
Predictor: Div. Yield						
Empirical OOS Statistics:		-0.35%			-0.10%	
PM OOS Statistics:	5.05	1.97%	2.32%	9.10	5.85%	5.95%
Overlap. Ann. Returns	Sample Begin: 1M 1872 Forecast Begin: 12M 1926			Sample Begin: 12M 1947 Forecast Begin: 12M 1967		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio						
Empirical OOS Statistics:		-1.41%			-0.15%	
PM OOS Statistics:	4.95	1.45%	2.86%	0.86	5.67%	5.82%
Predictor: Div. Yield						
Empirical OOS Statistics:		-0.85%			-0.03%	
PM OOS Statistics:	4.95	1.87%	2.72%	15.32	7.35%	7.38%

Panel B: ESV Prior

	Total Sample			**Postwar Sample**		
Annual Returns	Sample Begin: 1892 Forecast Begin: 1926			Sample Begin: 1947 Forecast Begin: 1967		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio						
Empirical OOS Statistics:		-0.55%			-0.87%	
ESV OOS Statistics:	0.38	-1.78%	-1.23%	0.28	-0.72%	0.15%
Predictor: Div. Yield						
Empirical OOS Statistics:		0.23%			-0.10%	
ESV OOS Statistics:	0.08	-4.23%	-4.46%	0.37	-0.85%	-0.75%
Overlap. Ann. Returns	Sample Begin: 12M 1914 Forecast Begin: 12M 1934			Sample Begin: 12M 1947 Forecast Begin: 12M 1967		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio						
Empirical OOS Statistics:		-3.12%			-0.15%	
ESV OOS Statistics:	0.38	-3.41%	-0.29%	0.26	-1.14%	-0.99%
Predictor: Div. Yield						
Empirical OOS Statistics:		-11.84%			-0.03%	
ESV OOS Statistics:	0.29	-6.57%	5.27%	0.26	-0.75%	-0.72%

Table 4.4: Results for the Habit Formation Term Structure Model-Based Priors

Reported are the results for the model-based priors from the HFTS model. The weight of the prior, λ, is estimated as explained in Section 4.3.2. The ΔR^2_{OS} stands for the difference between the R^2_{OS} of the predictor with and without model-based prior. A positive value indicates an improvement in the OOS forecasts through the prior. The data sample ends in December 2011.

Panel A: PM Prior

	Total Sample			**Postwar Sample**		
Quarterly Returns	Sample Begin: 1Q 1920 Forecast Begin: 1Q 1940			Sample Begin: 1Q 1947 Forecast Begin: 1Q 1967		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Short Term Yield						
Empirical OOS Statistics:		-0.19%			-3.50%	
PM OOS Statistics:	0.07	0.12%	0.31%	0.09	-1.24%	2.26%
Predictor: Long Term Yield						
Empirical OOS Statistics:		-1.72%			-3.44%	
PM OOS Statistics:	0.07	-0.34%	1.38%	0.19	-0.11%	3.33%
Predictor: Term Spread						
Empirical OOS Statistics:		0.03%			-3.61%	
PM OOS Statistics:	0.07	0.19%	0.16%	0.08	-1.09%	2.52%

Panel B: ESV Prior

	Postwar Sample		
Quarterly Returns	Sample Begin: 3Q 1947 Forecast Begin: 3Q 1967		
	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Short Term Yield			
Empirical OOS Statistics:		-3.08%	
ESV OOS Statistics:	0.03	-0.57%	2.51%
Predictor: Long Term Yield			
Empirical OOS Statistics:		-3.34%	
ESV OOS Statistics:	0.02	-0.71%	2.63%
Predictor: Term Spread			
Empirical OOS Statistics:		-3.03%	
ESV OOS Statistics:	0.03	-3.03%	0.00%

Table 4.3.

The PM prior leads to a substantial increase in OOS R^2 for the dividend-price ratio and the dividend yield. This performance over the total sample is not matched by the HF or the PT model. Further, the values for λ are very high. For the total sample, λ is around 5 for both predictors. For the postwar sample, the performance improvement is even more pronounced. The OOS R^2 is pushed from a negative value to over 7%. Hence, a real time investor could have profited significantly by forecasting the equity premium and imposing parameter restrictions derived from the LLR model through Bayesian priors. However, the performance of the ESV prior does not match the success of the PM prior. The ESV prior causes the OOS performance to deteriorate in all but two cases. A possible explanation is that when simulating data from the LRR for the ESV prior, four different time series are required as input variables: the two state variables, consumption growth, and dividend growth. This makes it more difficult for the model to generate data through which the ESV prior can improve the accuracy of the predictions.

4.5.4 Habit Formation Term Structure model-based priors

Model-based priors for the long term yield, the short term yield, and the term spread are derived from the HFTS model. Table 4.4 shows the results obtained from implementing these priors on predictive regressions with quarterly data.[10] For the PM prior the data sample starts in the 1st Quarter of 1920 and the OOS forecast begins in the 1st Quarter of 1940. The postwar sample starts in the 1st Quarter of 1947 for the PM prior. For the ESV prior, the data sample starts in the 3rd Quarter of 1947, since quarterly consumption data from FRED, which is needed to compute the log consumption growth process and generate data for the ESV prior, is unavailable before this.

The PM prior is again successful: it improves the OOS R^2 for all three predictive

[10]No overlapping returns are used as the HFTS model is simulated and matches empirical moments at a quarterly frequency.

variables and over both data samples. The λ values are very small, which implies that the majority of the weight is placed on the empirical data when estimating the parameters. The ESV prior can improve the OOS R^2 for both the short term and the long term yields by several percentage points despite the small values of λ. However, the OOS R^2 remain negative, as the OOS R^2 without imposing a prior are very low.

4.5.5 Sensitivity of the model-based prior's OOS performance to λ

The results above only consider a particular value of λ, namely the λ which maximizes the data density over the initial estimation period. An interesting question is whether the results are sensitive with respect to λ. There exist three cases. First, imposing a model-based prior helps to raise the OOS R^2 for small values of λ, but the performance deteriorates and the OOS R^2 falls below the no-prior level for a large λ. Second, any value of λ improves the forecasting performance. For this case, simply taking the theoretical parameter estimates implied by the asset pricing model will perform better than using empirical data for the estimation of the single variable predictive regression. This finding is similar in spirit to the result in Campbell and Thompson, which shows that using the theoretical parameter values implied by the steady state theory of Gordon (1962) and ignoring the empirical estimates of the parameters in the predictive regression can lead to more accurate forecasts. Third, the model-based prior fails to improve the OOS R^2 for any λ.

Figure 4.1 shows these three cases. For the PM prior from the HF model imposed on the dividend-price ratio over the postwar sample and forecasting annual returns, a λ between 0 and 3 raises the OOS R^2. For a λ greater than 3, the OOS performance deteriorates. When imposing the PM prior from the LRR model, the OOS R^2 converges to a level above 6%. Hence, the prior improves the predictive power for all values of λ. The opposite case is shown by the third graph, where the ESV prior

from the LLR model imposed on the dividend yield fails to improve the OOS R^2 for any value of λ.

Generally, the performance of the model-based priors is stable with respect to λ: a prior which improves the predictability of the equity premium does so for a range of λ values.

4.5.6 Is the model-based prior truly OOS?

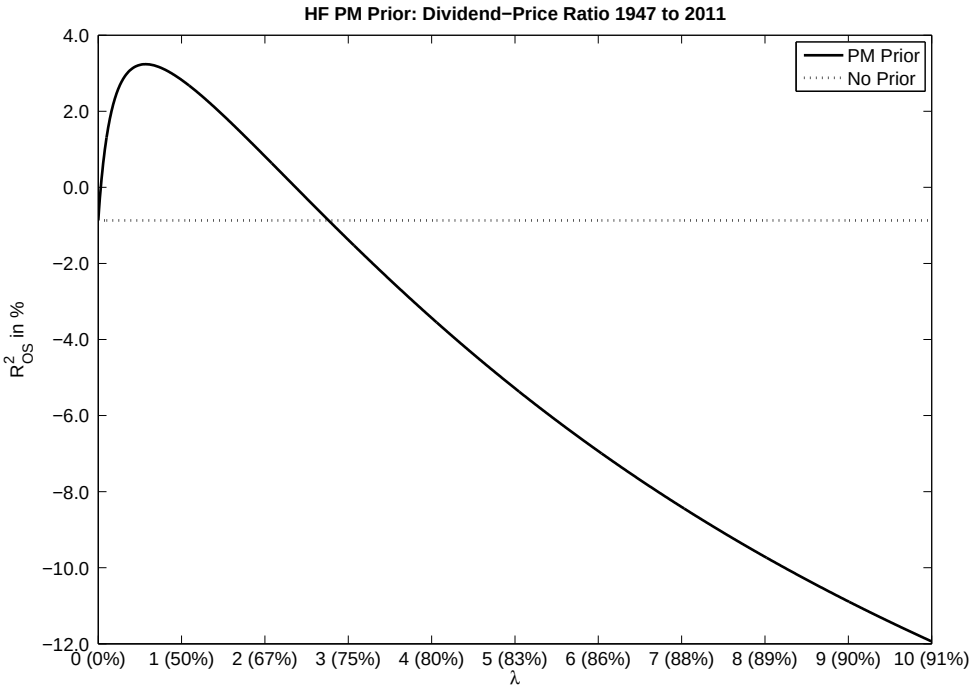
The results above show the forecast of the equity premium OOS. The equity premium is predicted recursively, and λ is estimated over an initial estimation period with no look-ahead bias. However, the parameters of the asset pricing models have been calibrated by the authors based on a data sample ending in the 1990s or 2000s, respectively, depending on the model.

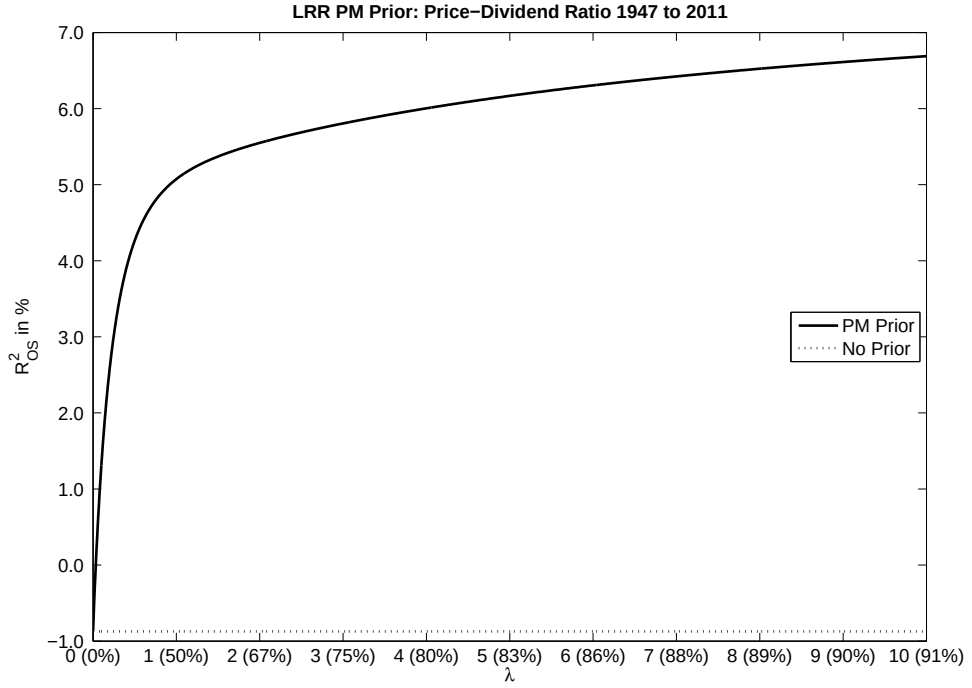
To assess whether the improvement in the predictive power holds up after the calibration data sample ended, I extend the initial estimation period up to the last month or quarter of the calibration data sample of a particular asset pricing model. For the model-based priors imposed at an annual frequency, overlapping annual returns are used to increase the number of OOS predictions.

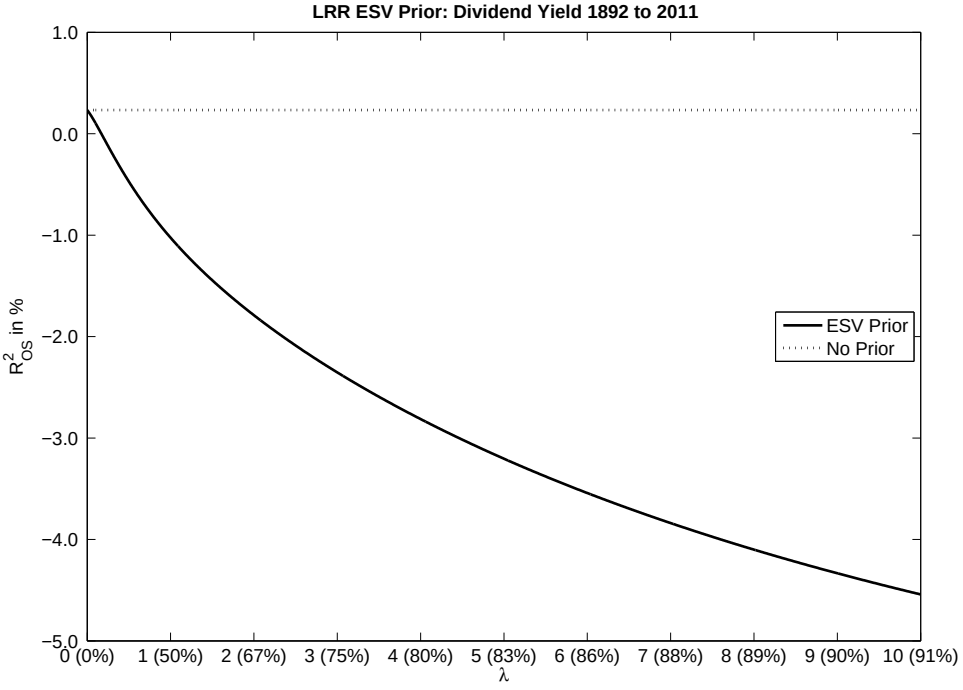
The results are reported in Table 4.5. For the HF model, the improvement in OOS predictability holds up, with the exception of the PM prior imposed on the dividend yield. The PM prior from the PT model is successful for both predictor variables (as described above, the ESV prior of the PT model is not used for overlapping annual returns due to the authors' model calibration). The priors from the LRR model are effective for the dividend yield but do not increase the OOS R^2 for the dividend-price ratio. The priors of the HFTS model are robust – the OOS R^2 is raised for all predictors and priors, except for the ESV prior used for the term spread. However, the ESV prior is also unsuccessful in improving the performance of the term spread in Table 4.4.

These results show that the model-based priors are successful when forecasting

Figure 4.1: Sensitivity of R^2_{OS} to λ







This figure depicts the sensitivity of R^2_{OS} to λ for three cases. The horizontal axis shows the value of λ and, in parentheses, the proportion of the simulated data relative to the total data (empirical and simulated data combined) used to estimate the parameters of the single variable predictive regression.

the equity premium over a sample period which does not contain any data used to calibrate the asset pricing model. However, the caveat is that the OOS periods are relatively short and based on a small number of independent predictions.

4.5.7 The model-based prior versus the Minnesota prior

The model-based priors impose economic restrictions on the parameter estimates of the single variable predictive regressions. An alternative is the Minnesota prior from Doan et al. (1984), which is purely statistical and not based on the economic theory implied by an economic model. The general purpose of the Minnesota prior is to reduce overfitting of VARs used in macroeconomics by shrinking the parameter estimates towards a random walk specification. I adapt the Minnesota prior such that it shrinks the estimate of β_0 towards the historical average and the estimate of β_1 towards zero – my version of the Minnesota prior shrinks the parameters towards the historical average return model. Comparing the forecast performance of the model-based prior and the Minnesota prior, helps to answer the question whether the economic theory derived from the asset pricing models adds value when being incorporated in a prior: the model-based priors should outperform the Minnesota prior if this is true.

My implementation of the Minnesota prior assumes the following prior distribution:

$$\beta \sim N\left(\underline{\beta},\ \underline{H}\right), \text{ where } \underline{H} = \begin{bmatrix} \phi_0 s_y^2 & 0 \\ 0 & \phi_1 \end{bmatrix} \text{ and } \underline{\beta} = \begin{pmatrix} \mu_y \\ 0 \end{pmatrix}. \tag{4.5.2}$$

The parameter μ_y is the mean and s_y^2 the variance of the equity premium over the estimation period. The prior tightness parameters for β_0 and β_1 are ϕ_0 and ϕ_1, respectively. The posterior distribution is then:

$$\bar{\beta} = \overline{H}[\underline{H}\underline{\beta} + (\hat{\sigma}^{-2}X)'Y], \text{ where } \overline{H} = [\underline{H}^{-1} + (\hat{\sigma}^{-2}X'X)]^{-1}. \tag{4.5.3}$$

Table 4.5: Results for the Post Calibration Period

Reported are the results for the forecast begin set such that the OOS forecasting period only contains data not used by the authors to calibrate the asset pricing model. The weight of the prior, λ, is estimated as explained in Section 4.3.2. The ΔR^2_{OS} stands for the difference between the R^2_{OS} of the predictor with and without model-based prior. A positive value indicates an improvement in the OOS forecasts through the prior. The postwar sample is used in all cases and the data end in December 2011. For the overlapping annual returns, the forecasting procedure is repeated in monthly steps, but the forecasts are at annual horizons.

Panel A: Habit Formation

Overlapping Annual Returns (Forecast Begin: December 1996)

	PM Prior			ESV Prior		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio						
Empirical OOS Statistics:		-8.62%			-8.62%	
PM OOS Statistics:	0.87	-5.58%	3.04%	0.22	-0.80%	7.70%
Predictor: Div. Yield						
Empirical OOS Statistics:		-0.71%			-0.71%	
PM OOS Statistics:	1.77	-3.05%	-1.92%	0.25	1.32%	2.03%

Panel B: Prospect Theory

Overlapping Annual Returns (Forecast Begin: December 1996)

	PM Prior		
	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio			
Empirical OOS Statistics:		-8.62%	
PM OOS Statistics:	0.11	-1.21%	7.41%
Predictor: Div. Yield Ratio			
Empirical OOS Statistics:		-0.71%	
PM OOS Statistics:	0.16	1.94%	2.65%

Panel C: Long Run Risk

Overlapping Annual Returns (Forecast Begin: December 1999)

	PM Prior			ESV Prior		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Div.-Price Ratio						
Empirical OOS Statistics:		16.36%			16.36%	
PM OOS Statistics:	5.28	10.10%	-6.35%	0.07	16.19%	-0.17%
Predictor: Div. Yield						
Empirical OOS Statistics:		12.35%			12.35%	
PM OOS Statistics:	5.28	12.85%	0.50%	0.08	12.54%	0.19%

Panel D: Habit Formation Term Structure

Quarterly Returns (Forecast Begin: 4Q 2004)

	PM Prior			ESV Prior		
	λ	R^2_{OS}	ΔR^2_{OS}	λ	R^2_{OS}	ΔR^2_{OS}
Predictor: Short Term Yield						
Empirical OOS Statistics:		-3.98%			-4.21%	
PM OOS Statistics:	0.28	-2.44%	1.54%	0.02	-3.69%	0.52%
Predictor: Long Term Yield						
Empirical OOS Statistics:		-1.04%			-1.12%	
PM OOS Statistics:	0.41	-0.28%	-0.10%	0.02	-0.89%	0.23%
Predictor: Term Spread						
Empirical OOS Statistics:		-3.93%			-3.94%	
PM OOS Statistics:	0.50	-2.10%	1.83%	0.02	-4.31%	-0.37%

To estimate $\hat{\sigma}^{-2}$, the OLS residuals of the single variable predictive regression, updated for each forecast, are used.

I set $\phi_0 = 1000$ such that the prior on β_0 is very diffuse – a diffuse prior on the constant of the model is common in applications of the Minnesota prior. For ϕ_1, I specify a Gamma distribution and average the posterior estimate $\bar{\beta}$ over the possible values of ϕ_1, i.e. I specify a hierarchical prior. This reflects the uncertainty of a real time investor when choosing ϕ_1. The Gamma distribution is used since the value of ϕ_1 has to be positive.

The results of a comparison between the two priors are reported in Table 4.6. For the model-based priors, the better performing prior, PM or ESV, for the respective period (total or postwar sample) and predictor of each asset pricing model is selected. The three different versions of the Minnesota prior correspond to changes in the scale parameter of the Gamma distribution. Version 1 corresponds to a Gamma distribution for ϕ_1 with a shape parameter of 2 and a scale parameter of 0.05. Versions 2 and 3 also use a shape parameter of 2, but the scale parameters are 0.01 and 0.001, respectively. These Gamma distributions lead to an average ϕ_1 which is line with the tightness parameters used in Doan et al. (1984). Adjusting the magnitude of the shape and scale parameter downwards changes the OOS R^2 only marginally. When the shape and scale parameters are increased, the performance of the Minnesota prior deteriorates.

For the dividend-price ratio and the dividend yield, the model-based prior raises the OOS R^2 to a higher level than the Minnesota prior in the majority of cases. This implies that the economic theory of the asset pricing models is valuable for the purpose of forecasting the equity premium. A purely statistical prior like the Minnesota prior does not achieve the same performance. For the yield related variables, the Minnesota prior generally fares slightly better than the model-based priors from the HFTS model. Hence, the economic insights from the HFTS model do not appear to result in a prior which is more useful than a purely statistical method for addressing

the problem of overfitting.

4.6 How the Model-Based Prior Improves Predictive Accuracy

To gain further insights into the working of the model-based prior, I analyse the two ways in which the model-based prior improves the OOS forecast performance of the single variable predictive regression in more detail. As mentioned in Section 4.2, these two ways are: the reduction in the overfitting of the parameter estimates and bringing the estimates of β_0 and β_1 closer to their in-sample counterparts during periods when the fundamental relationship between the predictor and the equity premium, which holds over the total sample, fails to hold temporarily.

Figure 4.2 depicts an example of a prior improving the OOS forecasts of the equity premium by smoothing the estimates of β and bringing them closer to their in-sample estimates. The figure shows the in-sample estimates over the total post-war sample, the empirical OOS estimates, and the OOS estimates when imposing the PM prior from the LRR model. The empirical estimates are more volatile since the overfitted parameter estimates pick up noise. Imposing the model-based prior helps to reduce the impact of noise on the parameter estimates: the posterior estimates are smoother when plotted over time. Further, the graph shows that the PM prior pushes the OOS estimates closer to the in-sample estimates for most of the OOS period. The empirical OOS estimates are far higher than their in-sample counterparts up to 1998, but the PM prior pushes the OOS estimates down. This also has the effect that for the end of the sample, from 1998 to 2011, the posterior estimates of the PM prior are lower than the in-sample estimates.

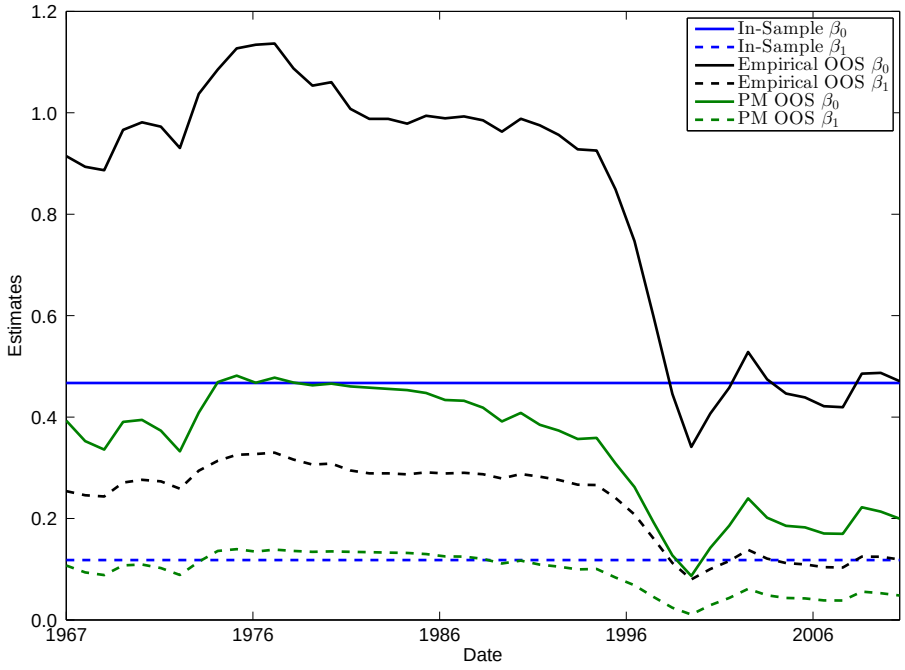
Table 4.7 reports the parameter estimates and the standard deviations over the total sample. The second and third column in the table show the in-sample estimates of the empirical data, of the data simulated from the LRR model through the Monte

Table 4.6: Comparison of the Model-Based Prior and the Minnesota Prior

Reported is the increase or decrease in R^2_{OS} when either imposing the model-based prior or the Minnesota prior. The baseline is the R^2_{OS} when no prior is imposed. For the model-based priors, the better performing prior, PM or ESV, for the respective period (total or postwar sample) and predictor of each model is selected. The initial estimation period is at least 20 years, but the OOS forecast does not start before 1926 for annual returns or December 1926 for overlapping annual returns. Version 1 of the Minnesota prior corresponds to the implementation of the prior with a Gamma distribution for ϕ_1 with a shape parameter of 2 and a scale parameter of 0.05. Versions 2 and 3 also use a shape parameter of 2, but the scale parameters are 0.01 and 0.001, respectively.

	Model-Based Prior ΔR^2_{OS}				**Minnesota Prior ΔR^2_{OS}**		
	HF	PT	LRR	HFTS	Version 1	Version 2	Version 3
Predictor: Div.-Price Ratio							
Annual Returns 1872-2011	0.01%	0.32%	1.86%		0.15%	0.46%	0.76%
Overlapping Annual Returns 12M 1872-12M 2011	1.16%	1.26%	2.86%		0.05%	0.20%	0.99%
Annual Returns 1947-2011	7.35%	6.52%	6.93%		1.72%	4.31%	4.48%
Overlappping Annual Returns 12M 1947-12M 2011	7.11%	5.94%	5.82%		0.16%	0.68%	3.37%
Predictor: Dividend Yield							
Annual Returns 1872-2011	1.02%	0.98%	2.32%		0.29%	0.70%	0.83%
Overlapping Annual Returns 12M 1872-12M 2011	1.59%	1.71%	2.72%		0.63%	0.85%	1.66%
Annual Returns 1947-2011	2.60%	5.06%	5.95%		1.05%	2.63%	2.57%
Overlapping Annual Returns 12M 1947-12M 2011	2.92%	5.82%	7.38%		0.22%	0.57%	2.31%
Predictor: Short Term Yield							
Quarterly Returns 1Q 1920-4Q 2011				0.31%	0.42%	0.34%	0.22%
Quarterly Returns 1Q 1947-4Q 2011				2.26%	4.01%	4.02%	3.56%
Predictor: Long Term Yield							
Quarterly Returns 1Q 1920-4Q 2011				1.38%	1.63%	1.73%	1.73%
Quarterly Returns 1Q 1947-4Q 2011				3.33%	3.42%	3.53%	3.42%
Predictor: Term Spread							
Quarterly Returns 1Q 1920-4Q 2011				0.16%	0.05%	0.01%	-0.01%
Quarterly Returns 1Q 1947-4Q 2011				2.52%	3.78%	3.64%	3.57%

Figure 4.2: Parameter Estimates over the Postwar Sample



This figure shows the parameter estimates over the postwar sample with the OOS period starting in 1967 for the dividend-price ratio and annual returns. The PM prior is from the LRR model and imposed with $\lambda = 1.9$.

Table 4.7: Parameter Estimates for the Postwar Sample

Reported are the in-sample parameter estimates of β_0 and β_1 for the dividend-price ratio using the postwar sample with and without model-based prior. Further, the standard deviation of the parameter estimates over the OOS period with forecast begin in 1967 using annual returns are shown. The estimates of the simulated data refer to the data generated by feeding empirical consumption growth into the HF model. The Monte Carlo simulation for the LLR model is described in Section 4.4.2.1.

	In-Sample Estimates		**OOS Statistics**	
	β_0	β_1	Std. Dev. of β_0 Estimates	Std. Dev. of β_1 Estimates
Habit Formation Model				
Empirical Data 1947-2011	0.467	0.118	0.258	0.082
Simulated Data 1947-2011	0.378	0.116	0.083	0.028
ESV Prior 1947-2011 ($\lambda = 0.73$)	0.269	0.066	0.137	0.045
Long Run Risk Model				
Empirical Data 1947-2011	0.467	0.118	0.258	0.082
Monte Carlo Simulation	0.089	0.016	NA	NA
PM Prior 1947-2011 ($\lambda = 1.90$)	0.197	0.048	0.119	0.039

Carlo procedure described in Section 4.4.2.1, and of the combined data, i.e. when imposing the PM prior, with the weight on the simulated data being $\lambda = 1.9$. We see that the theoretical estimates implied by the LRR model, $\hat{\beta}_0 = 0.089$ and $\hat{\beta}_1 = 0.016$, are substantially lower than the empirical estimates, $\hat{\beta}_0 = 0.467$ and $\hat{\beta}_1 = 0.118$. Thus, when imposing the PM prior, the resulting posterior estimates, $\hat{\beta}_0 = 0.197$ and $\hat{\beta}_1 = 0.048$, are lower than the empirical estimates.

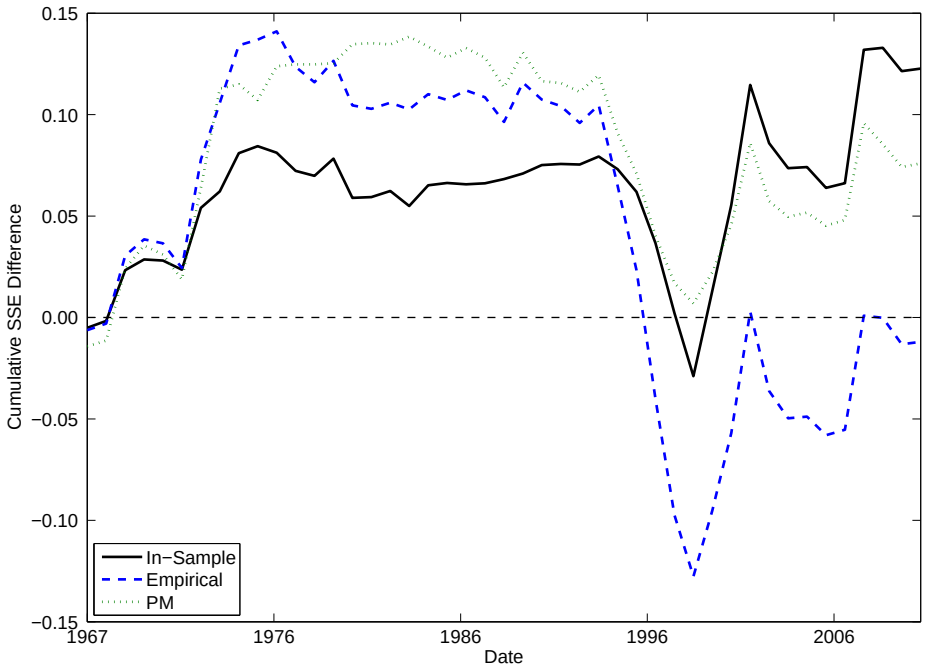
The fourth and fifth column report the standard deviations of the estimates over the OOS periods. The standard deviations of the empirical estimates are higher than those of the posterior estimates. This is a result of the smoothing effect of the model-based prior as shown in Figure 4.2.

Figure 4.3 shows why pushing the OOS estimates closer to the in-sample estimates for most of the sample period helps to increase the OOS R^2. The graph presents the difference between the cumulative sum of squared errors (SSE). I take the cumulative SSE of the historical average model and subtract the cumulative in-sample regression residuals, the cumulative SSE of the OOS predictive regression estimated only with empirical data, and the cumulative SSE of the OOS predic-

tive regression with the LRR PM prior imposed with $\lambda = 1.9$, respectively. A model which outperforms the historical average model, has a positive difference in cumulative SSE. We can see that the in-sample regression underperforms the OOS regression up to around 1997, but performs better at the end of our sample period, because the empirical model loses its previous gains over a window of four years. This is because during the Dot-com bubble the dividend-price ratios were very low and the empirical data up to this point implied a strong relationship between a low dividend-price ratio and a subsequent low equity premium. Hence, during the years leading up to the bursting of the Dot-com bubble, the empirical model predicts a negative equity premium. However, the negative equity premium did not materialise until the year 2000, which leads to the drop in the difference between the cumulative SSE for the predictive regression estimated only with empirical data. Since the in-sample coefficients are estimated over the total sample, the estimator can trade-off current losses and future gains. The in-sample estimates tolerate larger squared errors around 1975 in order to limit the squared errors in the late 1990s. Since the in-sample estimates include this period of low predictability in the late 1990s, they are smaller than the OOS estimates up to 1998. They have a forward-looking bias. The PM prior helps to avoid the pitfall of the Dot-com bubble and brings the difference in cumulative SSE curve closer to the in-sample curve. The PM prior achieves this as the LRR model implies that the dividend-price ratio has weak predictive power. Hence, the PM prior helps to anticipate periods of weak predictability, which leads to a reduction in forecasting errors.

The ESV prior of the HF model implies lower predictability and has the same effect as the PM prior from the LRR model. Table 4.7 reports the estimates of β with and without the ESV prior from the HF model. Further, the estimates of the simulated data generated by feeding empirical consumption growth to the HF model are also provided. The ESV prior of the HF model pushes the parameter estimates down, since the model implies less predictive power of the dividend-price ratio than

Figure 4.3: Cumulative Sum of Squared Errors over the Postwar Sample



This figure shows the difference between the cumulative sum of squared errors (SSE) of the historical average model and either the cumulative in-sample regression residuals, the cumulative SSE of the OOS predictive regression estimated only with empirical data, or the cumulative SSE of the OOS predictive regression with the LRR PM prior imposed with $\lambda = 1.9$. A model which outperforms the historical average model has a positive difference in cumulative SSE. The predictor is the dividend-price ratio, and annual returns are used.

found in the empirical data. Hence, the ESV prior of the HF model also reduces the forecasting errors during periods of low predictability. The posterior estimates are even lower than the theoretical estimates implied by the model. For small values of λ, the estimates can be pushed down below or pushed up above the model implied estimates. However, as $\lambda \to \infty$, the posterior estimates converge to the estimates implied by the asset pricing model.

4.7 Conclusion

In this chapter, I impose restrictions on the parameter estimates of typical single variable predictive regressions used in the equity premium prediction literature. My parameter restrictions are based on a well specified Bayesian framework which has been inspired by the work of Del Negro and Schorfheide (2004). I adapt their framework to derive priors from consumption-based asset pricing models. These priors are termed model-based priors and are derived from from four asset pricing models: Habit Formation (see Campbell and Cochrane (1999)), Prospect Theory (see Barberis et al. (2001)), Long Run Risk (see Bansal and Yaron (2004)), and Habit Formation Term Structure (see Wachter (2006)).

I find that implementing the priors on the single variable predictive regressions can lead to increases in the OOS R^2 of up to several percentage points. The model-based priors perform particularly well for the postwar period when using the dividend-price ratio or the dividend yield as predictors. The prior tightness is determined by the hyperparameter λ, which is estimated from the data. The estimation of λ is only based on data which are available before the beginning of the OOS forecast. Hence, a real time investor could implement the estimation procedure described in this chapter and achieve substantially higher expected returns from the higher OOS R^2.

I also compare the performance of the model-based priors with a purely statistical

one, namely the Minnesota prior, which does not incorporate any economic theory derived from an asset pricing model. In most cases, the model-based priors raise the OOS R^2 by more than the Minnesota prior: this implies that the economic theory of the asset pricing models is valuable for the purpose of forecasting the equity premium. Further, the results are also robust when restricting the OOS period to a subsample which does not include any data used for the calibration of the asset pricing model by the respective authors.

Appendix

4.A Model-Based Priors

4.A.1 The prior and posterior distributions

The model-based prior is based on the DSGE-VAR approach of Del Negro and Schorfheide (2004). The likelihood function for the predictive regression based on the total empirical data sample of size T, denoted by (Y, X), is

$$p(Y, X \mid \beta, \sigma^2) \propto (\sigma^2)^{-\frac{T}{2}} exp\left\{-\frac{1}{2}\sigma^{-2}(Y'Y - \beta' X'Y - Y'X\beta + \beta' X'X\beta)\right\}. \tag{4.A.1}$$

The likelihood function for the predictive regression based on data simulated from the asset pricing model (Y^*, X^*) is accordingly

$$p(Y^*(\theta), X^*(\theta) \mid \beta, \sigma^2) \propto (\sigma^2)^{-\frac{\lambda T}{2}} \times exp\left\{-\frac{1}{2}\sigma^{-2}(Y^{*\prime}Y^* - \beta' X^{*\prime}Y^* - Y^{*\prime}X^*\beta + \beta' X^{*\prime}X^*\beta)\right\}. \tag{4.A.2}$$

The asset pricing model parameters are denoted by θ, and λ is the hyperparameter which determines the weight of the simulated data in the posterior. These individual likelihoods are combined to form the joint likelihood:

$$p(Y^*(\theta), Y, X^*(\theta), X \mid \beta, \sigma^2) = p(Y^*(\theta), X^*(\theta) \mid \beta, \sigma^2) p(Y, X \mid \beta, \sigma^2), \tag{4.A.3}$$

where $p(Y^*(\theta), X^*(\theta) \mid \beta, \sigma^2)$ can be interpreted as a prior for the parameters of the predictive regression. I specify this prior in two different ways: the PM prior and the ESV prior.

Through the use of scaled population moments, we can rewrite (4.A.2) as

$$\begin{aligned} p(\beta, \sigma^2 \mid \theta) =& c_{PM}(\theta)^{-1}(\sigma^2)^{-\frac{\lambda T+2}{2}} \\ & \times exp\left\{-\frac{1}{2}\lambda T\sigma^{-2}(\Gamma^*_{yy}(\theta) - \beta'\Gamma^*_{xy}(\theta) - \Gamma^*_{yx}(\theta)\beta + \beta'\Gamma^*_{xx}(\theta)\beta)\right\}. \end{aligned} \tag{4.A.4}$$

We incorporate an improper prior $p(\beta, \sigma^2) \propto \sigma^{-2}$ together with $c_{PM}(\theta)$, a normalisation factor which ensures that the density integrates to one. The form of $c_{PM}(\theta)$ is given in Appendix 4.A.2. The PM prior distribution of the predictive regression parameters is of the Inverted-Wishart Normal form:

$$\sigma^2_{PM} \mid \theta \backsim IW(\lambda T\sigma^{2*}_{PM}, \lambda T - 2, 1) \text{ and} \tag{4.A.5}$$

$$\beta_{PM} \mid \sigma^2_{PM}, \theta \backsim N(\beta^*_{PM}(\theta), \sigma^2_{PM} \otimes (\lambda T\Gamma^*_{xx}(\theta))^{-1}), \tag{4.A.6}$$

where

$$\beta^*_{PM}(\theta) = \Gamma^*_{xx}(\theta)^{-1}\Gamma^*_{xy}(\theta) \text{ and} \tag{4.A.7}$$

$$\sigma^{2*}_{PM}(\theta) = \Gamma^*_{yy}(\theta) - \Gamma^*_{yx}(\theta)\Gamma^*_{xx}(\theta)^{-1}\Gamma^*_{xy}(\theta). \tag{4.A.8}$$

The ESV prior takes the form

$$\begin{aligned} p(\beta, \sigma^2 \mid \theta) =& c_{ESV}(\theta)^{-1}(\sigma^2)^{-\frac{\lambda T+2}{2}} \\ & \times exp\left\{-\frac{1}{2}\lambda\sigma^{-2}(\bar{Y}'\bar{Y}(\theta) - \beta'\bar{X}'\bar{Y}(\theta) - \bar{Y}'\bar{X}(\theta)\beta + \beta'\bar{X}'\bar{X}(\theta)\beta)\right\}. \end{aligned} \tag{4.A.9}$$

Again, we add an improper prior $p(\beta, \sigma^2) \propto \sigma^{-2}$ together with the normalisation factor $c_{ESV}(\theta)$. The form of $c_{ESV}(\theta)$ is given in Appendix 4.A.2. The ESV prior

distribution of the predictive regression parameters is

$$\sigma^2_{ESV} \mid \theta \backsim IW(\lambda T \sigma^{2*}_{ESV}, \lambda T - 2, 1) \text{ and} \tag{4.A.10}$$

$$\beta_{ESV} \mid \sigma^2_{ESV}, \theta \backsim N(\beta^*_{ESV}(\theta), \sigma^2_{ESV} \otimes (\lambda \bar{X}' \bar{X})^{-1}), \tag{4.A.11}$$

where

$$\beta^*_{ESV}(\theta) = (\bar{X}' \bar{X}(\theta))^{-1} \bar{X}' \bar{Y}(\theta) \text{ and} \tag{4.A.12}$$

$$\sigma^{2*}_{ESV}(\theta) = \frac{1}{T}(\bar{Y}' \bar{Y}(\theta) - \bar{Y}' \bar{X}(\theta)(\bar{X}' \bar{X}(\theta))^{-1} \bar{X}' \bar{Y}(\theta)). \tag{4.A.13}$$

For the PM prior, the maximum-likelihood estimates of β and σ^2 based on simulated and empirical data take the following form

$$\tilde{\beta}_{PM}(\theta) = (\lambda T \Gamma^*_{xx}(\theta) + X'X)^{-1}(\lambda T \Gamma^*_{xy}(\theta) + X'Y) \text{ and} \tag{4.A.14}$$

$$\begin{aligned} \tilde{\sigma}^2_{PM}(\theta) =& \frac{1}{(\lambda+1)T}[(\lambda T \Gamma^*_{yy}(\theta) + Y'Y) \\ & - (\lambda T \Gamma^*_{yx}(\theta) + Y'X)(\lambda T \Gamma^*_{xx}(\theta) + X'X)^{-1}(\lambda T \Gamma^*_{xy}(\theta) + X'Y)]. \end{aligned} \tag{4.A.15}$$

The posterior distribution of the predictive regression parameters when implementing the PM prior is given by

$$\sigma^2_{PM} \mid Y, \theta \backsim IW((\lambda+1)T \tilde{\sigma}^2_{PM}(\theta), (\lambda+1)T - 2, 1) \text{ and} \tag{4.A.16}$$

$$\beta_{PM} \mid Y, \sigma^2_{PM}, \theta \backsim N(\tilde{\beta}_{PM}(\theta), \sigma^2_{PM} \otimes (\lambda T \Gamma^*_{xx}(\theta) + X'X)^{-1}). \tag{4.A.17}$$

For the ESV prior, the maximum-likelihood estimates of β and σ^2 are

$$\tilde{\beta}_{ESV}(\theta) = (\lambda \bar{X}' \bar{X}(\theta) + X'X)^{-1}(\lambda \bar{X}' \bar{Y}(\theta) + X'Y) \text{ and} \tag{4.A.18}$$

$$\begin{aligned} \tilde{\sigma}^2_{ESV}(\theta) =& \frac{1}{(\lambda+1)T}[(\lambda \bar{Y}' \bar{Y}(\theta) + Y'Y) \\ & - (\lambda \bar{Y}' \bar{X}(\theta) + Y'X)(\lambda \bar{X}' \bar{X}(\theta) + X'X)^{-1}(\lambda \bar{X}' \bar{Y}(\theta) + X'Y)]. \end{aligned} \tag{4.A.19}$$

Hence, the posterior distribution takes the form

$$\sigma^2_{ESV} \mid Y, \theta \backsim IW((\lambda+1)T\tilde{\sigma}^2_{ESV}(\theta), (\lambda+1)T-2, 1) \text{ and} \tag{4.A.20}$$

$$\beta_{ESV} \mid Y, \sigma^2_{ESV}, \theta \backsim N(\tilde{\beta}_{ESV}(\theta), \sigma^2_{ESV} \otimes (\lambda\bar{X}'\bar{X}(\theta) + X'X)^{-1}). \tag{4.A.21}$$

Since I take θ as given by the authors of the asset pricing models – $p(\theta)$ in Del Negro and Schorfheide (2004) is set equal to one – drawing from the posterior distribution is trivial. For β, the modes are $\tilde{\beta}_{PM}(\theta)$ and $\tilde{\beta}_{ESV}(\theta)$, depending on the prior. For σ^2, the modes are $\tilde{\sigma}^2_{PM}$ and $\tilde{\sigma}^2_{ESV}$, respectively.

4.A.2 Normalisation factors

For the PM prior, the normalisation factor $c_{PM}(\theta)$ in (4.A.4) is

$$c_{PM}(\theta) = 2\pi \mid \lambda T\Gamma^*_{xx}(\theta) \mid^{-\frac{1}{2}} \mid \lambda T\sigma^{*2}_{PM}(\theta) \mid^{-\frac{\lambda T-2}{2}} 2^{\frac{\lambda T-2}{2}}\Gamma[(\lambda T-2)/2], \tag{4.A.22}$$

and for the ESV prior, the normalisation factor $c_{ESV}(\theta)$ in (4.A.9) takes the form

$$c_{ESV}(\theta) = 2\pi \mid \lambda\bar{X}'\bar{X}(\theta) \mid^{-\frac{1}{2}} \mid \lambda T\sigma^{*2}_{ESV}(\theta) \mid^{-\frac{\lambda T-2}{2}} 2^{\frac{\lambda T-2}{2}}\Gamma[(\lambda T-2)/2], \tag{4.A.23}$$

where $\Gamma(.)$ is the gamma function.

4.A.3 Data densities

For the PM prior, the data density $p^{PM}_{\lambda}(Y, X \mid \theta)$ is

$$\begin{aligned} p^{PM}_{\lambda}(Y, X \mid \theta) =& \frac{\mid \lambda T\Gamma^*_{xx}(\theta) + X'X \mid^{-\frac{1}{2}} \mid (\lambda+1)T\tilde{\sigma}^2_{PM}(\theta) \mid^{-\frac{(\lambda+1)T-2}{2}}}{\mid \lambda T\Gamma^*_{xx}(\theta) \mid^{-\frac{1}{2}} \mid \lambda T\sigma^{*2}_{PM}(\theta) \mid^{-\frac{\lambda T-2}{2}}} \\ & \times \frac{(2\pi)^{-\frac{T}{2}} 2^{\frac{((\lambda+1)T-2)}{2}}\Gamma[((\lambda+1)T-2)/2]}{2^{\frac{\lambda T-2}{2}}\Gamma[(\lambda T-2)/2]}, \end{aligned} \tag{4.A.24}$$

and for the ESV prior, the data density $p_{\lambda}^{ESV}(Y, X \mid \theta)$ takes the form

$$\begin{aligned} p_{\lambda}^{ESV}(Y, X \mid \theta) =& \frac{\mid \lambda \bar{X}'\bar{X}(\theta) + X'X \mid^{-\frac{1}{2}} \mid (\lambda+1)T\tilde{\sigma}_{ESV}^{2}(\theta) \mid^{-\frac{(\lambda+1)T-2}{2}}}{\mid \lambda \bar{X}'\bar{X}(\theta) \mid^{-\frac{1}{2}} \mid \lambda T \sigma_{ESV}^{*2}(\theta) \mid^{-\frac{\lambda T-2}{2}}} \\ & \times \frac{(2\pi)^{-\frac{T}{2}} 2^{\frac{((\lambda+1)T-2)}{2}} \Gamma[((\lambda+1)T-2)/2]}{2^{\frac{\lambda T-2}{2}} \Gamma[(\lambda T-2)/2]}. \end{aligned} \tag{4.A.25}$$

These forms are derived by using the normalization constants for the Inverted-Wishart Normal distribution as shown in Zellner (1971).

4.B Asset Pricing Models

4.B.1 By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior

Campbell and Cochrane (1999) use a standard representative-agent consumption-based asset pricing model but add a slow-moving habit to the basic power utility function. This slow-moving habit leads to a slowly time-varying risk aversion and an equity risk premium which is higher at business cycle troughs than at peaks.

The agents are identical and maximise their utility given by

$$E\left[\sum_{t=0}^{\infty} \delta^{t} \frac{(C_t - X_t)^{(1-\gamma)} - 1}{1-\gamma}\right], \tag{4.B.1}$$

where X_t is the level of habit, and δ is the time discount factor. A surplus consumption ratio $S_t \equiv (C_t - X_t)/C_t$ is defined – a small value of S_t indicates that the economy is in a bad state.

A process is specified for $s_t = ln(S_t)$ which ensures that C_t is always above X_t. This process is given by

$$s_{t+1} = (1-\phi)\bar{s} + \phi s_t + \Psi(s_t)(c_{t+1} - c_t - g), \tag{4.B.2}$$

with ϕ reflecting habit persistence, i.e. how quickly s_{t+1} returns to the steady state value $\bar{s}$. The function $\Psi(s_t)$ is specified as

$$\Psi(s_t) = \begin{cases} \frac{1}{\bar{S}}\sqrt{1-2(s_t-\bar{s})}-1, & s_t \leq s_{max} \\ 0 & s_t > s_{max} \end{cases}, \tag{4.B.3}$$

with the parameter s_{max} set equal to $\bar{s} + \frac{1}{2}(1-\bar{S}^2)$. The steady state value of s_t, i.e. $\bar{s}$, is given by $ln(\sigma\sqrt{\gamma/(1-\phi)})$. The evolution of s_{t+1} is based on consumption growth being an i.i.d. lognormal process:

$$\Delta c_{t+1} = g + v_{t+1}, \text{ where } v_{t+1} \overset{i.i.d.}{\sim} N(0, \sigma_v^2). \tag{4.B.4}$$

Stocks represent a claim to the consumption stream. The price-consumption ratio for a consumption claim satisfies

$$\frac{P_t}{C_t}(s_t) = E_t\left[M_{t+1}\frac{C_{t+1}}{C_t}\left[1+\frac{P_{t+1}}{C_{t+1}}(s_{t+1})\right]\right]. \tag{4.B.5}$$

The underlying assumption is that dividend growth is perfectly correlated with consumption growth in (4.B.4).[11]

The price-consumption ratio is correlated with the business cycles as it depends on s_t. The ratio is high at business cycle peaks and low at troughs. Due to the slowly time-varying risk aversion, the equity premium is also correlated with the business cycle, but this correlation is negative. Hence, the model generates an equity risk premium which is predictable by the price-consumption ratio. A high price-consumption ratio predicts a low equity premium.

I apply the Fixed-Point Method to solve for the price-consumption and the price-dividend ratio (Wachter (2005)).

[11]Results for a model specification which assumes imperfectly correlated consumption and dividend processes are given in the paper. However, in this chapter, I focus only on the baseline specification.

4.B.2 A Consumption-Based Model of the Term Structure of Interest Rates

Wachter's (2006) model is based on the HF model but differs in two ways. First, Wachter departs from the constant risk-free rate assumption by using a different specification for $\bar{S}$. Second, she adds an inflation process, and thus, a nominal term structure can be computed.

Wachter specifies $\Psi(s_t)$ in the same way as Campbell and Cochrane (1999). However, $\bar{s}$ which is set equal to $ln(\sigma\sqrt{\gamma/(1-\phi-b/\gamma)})$ differs due to the parameter b. Hence, the value s_{max}, defined as $\bar{s}+\frac{1}{2}(1-\bar{S}^2)$, is also different. In the HF model, b is set equal to zero, which implies a constant risk-free rate. In the HFTS model, the log risk-free rate is given by

$$r_{f,t+1} = -ln(\delta) + \gamma g - \frac{\gamma(1-\phi)-b}{2} + b(\bar{s}-s_t). \tag{4.B.6}$$

The parameter b is chosen to be positive. Thus, an increase in s_t leads to a lower risk-free rate.

Wachter denotes the exogenous price level Π_t, which implies that inflation is given by $ln(\Pi_{t+1}/\Pi_t) = \triangle\pi_{t+1}$. Thus, the nominal bond pricing equation is

$$P^{\$}_{n,t} = E_t\left[M_{t+1}\frac{\Pi_t}{\Pi_{t+1}}P^{\$}_{n-1,t+1}\right] = E_t\left[\delta\left(\frac{S_{t+1}}{S_t}\frac{C_{t+1}}{C_t}\right)^{-\gamma}\frac{\Pi_t}{\Pi_{t+1}}P^{\$}_{n-1,t+1}\right], \tag{4.B.7}$$

with $P^{\$}_{0,t} = 1$. M_{t+1} is the pricing kernel, and n denotes the time to maturity.

While consumption growth is specified as an i.i.d. lognormal process of the form

$$\Delta c_{t+1} = g + \sigma_c\epsilon_{c,t+1}, \text{ where } \epsilon_{c,t+1} \overset{i.i.d.}{\sim} N(0,1), \tag{4.B.8}$$

the inflation process is given by

$$\Delta\pi_{t+1} = \eta_0 + \eta Z_t + \sigma_\pi\epsilon_{\pi,t+1}, \tag{4.B.9}$$

where the expected inflation Z_t takes the form

$$Z_t = (\Delta\pi_{t+1} - \eta_0 - \sigma_\pi \epsilon_{\pi,t+1}). \tag{4.B.10}$$

Z_t is univariate and follows an AR(1) process:

$$Z_t = \mu + \Phi Z_{t-1} + \Sigma \epsilon_{\pi,t}. \tag{4.B.11}$$

The shocks to consumption growth and to inflation are imperfectly correlated. The correlation is denoted by ρ.

Due to the inflation process, the nominal bond price depends on two state variables, namely Z_t and s_t. This complicates solving (4.B.7) as there are two random shocks, $\epsilon_{c,t+1}$ and $\epsilon_{\pi,t+1}$. However, through conditioning on $\sigma_c \epsilon_{c,t+1}$ and the Law of Iterated Expectations, the nominal bond price equals

$$P^{\$}_{n,t} = F^{\$}_n(s_t) exp\{A_n + B_n Z_t\}, \tag{4.B.12}$$

where

$$\begin{aligned} F^{\$}_n(s_t) =& E_t \left[M_{t+1} exp\{\xi_n \sigma_c \epsilon_{c,t+1}\} F^{\$}_{n-1}(s_{t+1}) \right] \\ =& E_t [exp\{ln(\delta) - \gamma g - \gamma(1-\phi)(\bar{s} - s_t) \\ & + [\xi_n - \gamma(\Psi(s_t)+1)]\sigma_c \epsilon_{c,t+1}\} F^{\$}_{n-1}(s_{t+1})], \end{aligned} \tag{4.B.13}$$

and does not depend on $\epsilon_{\pi,t+1}$. Furthermore, A_n and B_n are scalars and defined recursively. The exact forms of A_n and B_n as well as the derivation of (4.B.12) can be found in Wachter (2006).

The yields are correlated with the business cycles as they depend on s_t. The correlation is negative due to the positive parameter b in equation (4.B.6). As the equity premium is also negatively correlated with the business cycles, low yields predict a low equity premium in the HFTS model.

I solve the model by using standard numerical integration methods.

4.B.3 Prospect Theory and Asset Prices

In the model of Barberis et al. (2001), the agent does not only derive utility from consumption but also from fluctuations of her or his financial wealth. There are two important aspects in the way financial wealth fluctuations affect the utility of an economic agent. First, the agent is loss averse. Second, the degree of loss aversion depends on prior investment outcomes. Prior gains lead to less loss aversion, and prior losses lead to more loss aversion. Hence, the risk aversion of the agent varies over time, as in the HF model.

Aggregate consumption growth and dividend growth follow the i.i.d. lognormal processes given by

$$\Delta c_{t+1} = g_c + \sigma_c \epsilon_{c,t+1}, \text{ where } \epsilon_{c,t+1} \overset{i.i.d.}{\sim} N(0,1) \tag{4.B.14}$$

and

$$\Delta d_{t+1} = g_d + \sigma_d \epsilon_{d,t+1}, \text{ where } \epsilon_{d,t+1} \overset{i.i.d.}{\sim} N(0,1), \tag{4.B.15}$$

with $corr(\epsilon_{c,t+1}, \epsilon_{d,t+1}) = \omega$.[12]

The agent's maximization problem is set up as

$$E\left[\sum_{t=0}^{\infty}\left(\delta^t \frac{C_t^{1-\gamma}}{1-\gamma} + b_0 \bar{C}_t^{-\gamma}\delta^{t+1} v(X_{t+1}, S_t, z_t)\right)\right]. \tag{4.B.16}$$

The second term captures that the agent cares about fluctuations in financial wealth. The variable X_{t+1} denotes the change of the financial wealth between time t and $t+1$ and is defined as

$$X_{t+1} \equiv S_t R_{t+1} - S_t R_{f,t}. \tag{4.B.17}$$

The variable S_t measures the value of the agent's risky assets at time t. The variable

[12] Barberis et al. (2001) consider two different specifications: Economy I, in which dividends equal consumption, and Economy II, in which consumption and dividends follow separate processes. The simulated moments of Economy II are much more successful in matching the empirical moments, and hence, I do not consider Economy I.

z_t accounts for prior gains and losses up to time t and is defined as Z_t/S_t, where Z_t is a historical benchmark level for the value of the risky asset. If z_t is smaller than one, the agent has prior gains, if z_t is greater than one, the agent faces prior losses. The discount factor is δ, and $b_0\bar{C}_t^{-\gamma}$ is a scaling term. The form of the utility function $v(.)$ is different conditional on prior gains or prior losses. More details can be found in the appendix.

The dynamics of z_t are given by the process

$$z_{t+1} = \eta\left(z_t \frac{\bar{R}}{R_{t+1}}\right) + (1-\eta). \tag{4.B.18}$$

The benchmark level Z_t reacts sluggishly to changes in the stock price. When S_t increases, Z_t should increase by less in order to allow for prior gains. The sluggishness is determined by the parameter $\eta \in [0,1]$. The closer η is to one, the more sluggish the benchmark level becomes. The parameter $\bar{R}$ is chosen such that the median value of z_t is around one.

The price-dividend ratio is assumed to be a function of the state variable z_t:

$$f_t \equiv P_t/D_t = f(z_t). \tag{4.B.19}$$

The real stock returns are thus given as

$$R_{t+1} = \frac{1+f(z_{t+1})}{f(z_t)} e^{g_d + \sigma_d \epsilon_{d,t+1}}. \tag{4.B.20}$$

Barberis et al. (2001) show that the equilibrium is characterized by a constant real risk-free rate,

$$R_f = \delta^{-1} e^{\gamma g_c - \gamma^2 \sigma_c^2/2}, \tag{4.B.21}$$

and a price-dividend ratio given by

$$\begin{aligned} 1 =& \delta e^{g_d - \gamma g_c + \gamma^2 \sigma_c^2 (1-\omega^2)/2} E_t \left[\frac{1 + f(z_{t+1})}{f(z_t)} e^{(\sigma_d - \gamma \omega \sigma_c)\epsilon_{d,t+1}} \right] \\ & + b_0 \delta E_t \left[\hat{v} \left(\frac{1 + f(z_{t+1})}{f(z_t)} e^{g_d + \sigma_d \epsilon_{d,t+1}}, z_t \right) \right], \end{aligned} \tag{4.B.22}$$

where the utility function $\hat{v}(R_{t+1}, z_t)$ is equal to $v(R_{t+1}, S_t, z_t)/S_t$.

As the HF model, the PT model is able to generate predictability in returns. A decrease in z_t leads to both, a higher price-dividend ratio and a less risk averse investor. When z_t falls, subsequent returns will be lower as less compensation for risk is required. Hence, a high price-dividend ratio predicts a low equity risk premium.

I solve the model by following the process laid out by Barberis et al. (2001).

4.B.4 Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles

Bansal and Yaron (2004) propose a solution to the equity premium puzzle through a consumption-based asset pricing model with Epstein and Zin (1989) preferences. Their model differs from other consumption-based asset pricing models in two ways. First, they include a small persistent expected growth rate component in the consumption and dividend growth rate processes. This causes consumption growth and the return on the market portfolio to covary positively, and hence, the economic agents require a higher risk premium. Second, they allow for time-varying volatility, which accounts for fluctuating economic uncertainty, in both processes: this additional source of systematic risk increases the risk premium further.

The asset pricing restriction for the real return on the market portfolio $R_{m,t+1}$, according to the Epstein and Zin (1989) preferences, is

$$E_t \left[\delta^\theta G_{c,t+1}^{-\frac{\theta}{\psi}} R_{c,t+1}^{-(1-\theta)} R_{m,t+1} \right] = 1, \tag{4.B.23}$$

where $G_{c,t+1}$ is the aggregate gross growth rate of consumption, $R_{c,t+1}$ denotes the real return on an asset which pays aggregate consumption as dividends, and δ is the time discount factor. The parameter θ is defined as $(1-\gamma)/(1-\frac{1}{\psi})$, where γ is the risk aversion parameter, and ψ accounts for the Intertemporal Elasticity of Substitution. To derive the real returns, the authors use the standard approximation of Campbell and Shiller (1988). The real log return for the claim to aggregate consumption is

$$r_{c,t+1} = \kappa_0 + \kappa_1 z_{t+1} - z_t + g_{c,t+1}, \tag{4.B.24}$$

where $g_{c,t+1}$ is the log consumption growth, and z_t denotes the log price-consumption ratio. The specification for the real log return on the market portfolio is

$$r_{m,t+1} = \kappa_{0,m} + \kappa_{1,m} z_{m,t+1} - z_{m,t} + g_{d,t+1}, \tag{4.B.25}$$

where $g_{d,t+1}$ is the log dividend growth rate, and $z_{m,t}$ denotes the log price-dividend ratio. The κ's and κ_m's are constants which depend on the average level of the price-consumption ratio or the price-dividend ratio, respectively.[13]

The dynamics of log consumption growth and log dividend growth, which incorporate a small persistent predictable component x_t, the long run risk component, and a time-varying volatility σ_t, reflecting fluctuating economic uncertainty, are

$$\begin{aligned} x_{t+1} =& \rho x_t + \varphi_e \sigma_t e_{t+1} \\ g_{c,t+1} =& \mu_c + x_t + \sigma_t \eta_{t+1} \\ g_{d,t+1} =& \mu_d + \phi x_t + \varphi_d \sigma_t u_{t+1} \\ \sigma^2_{t+1} =& \sigma^2 + v_1(\sigma^2_t - \sigma^2) + \sigma_w w_{t+1}, \end{aligned} \tag{4.B.26}$$

with e_{t+1}, u_{t+1}, η_{t+1}, and w_{t+1} having i.i.d. standard Normal distributions. The

[13] Bansal, Kiku, and Yaron (2010) define κ_1 as $exp(\bar{z})/(1+exp(\bar{z}))$ and set κ_0 equal to $ln(1+exp(\bar{z})) - \kappa_1 \bar{z}$, where $\bar{z}$ is the mean log price-consumption ratio. Accordingly, $\kappa_{1,m}$ is defined by $exp(\bar{z}_m)/(1+exp(\bar{z}_m))$, and $\kappa_{0,m}$ is set equal to $ln(1+exp(\bar{z}_m)) - \kappa_{1,m}\bar{z}_m$, with $\bar{z}_m$ being the mean log price-dividend ratio.

state variables, which determine the price-consumption and price-dividend ratio, are x_t and σ_t. The solutions for z_t and $z_{m,t}$ are

$$\begin{aligned} z_t =& A_0 + A_1 x_t + A_2 \sigma_t^2 \\ z_{m,t} =& A_{0,m} + A_{1,m} x_t + A_{2,m} \sigma_t^2. \end{aligned} \tag{4.B.27}$$

The derivation of the A and A_m can be found in Bansal and Yaron (2004) and also in Bansal, Kiku, and Yaron (2010).

The model generates excess returns which are predictable by the price-dividend ratio. Equation (4.B.25) shows that the lagged price-dividend ratio has a negative impact on future returns. Hence, the relation implied between the price-dividend ratio and future returns is the same as in the HF and PT model.

4.C Additional Tables

Table 4.8: Comparison of the Model-Based Prior and the Campbell and Thompson Restrictions

Reported is the increase or decrease in R^2_{OS} when either imposing the model-based priors or the parameter restrictions from Campbell and Thompson (2008). For the model-based priors, the better performing prior, PM or ESV, for the respective period (total or postwar sample) and predictor of each model is selected. Bounded Slope means that the coefficient of the predictor must be of the same sign as the in-sample estimate, otherwise it is set to zero. The restriction Bounded Intercept/Slope ensures that the intercept is positive, otherwise it is set to zero, and the slope is between zero and one, otherwise it is truncated at these values. Fixed Coefficients stands for setting the slope equal to zero and the coefficient equal to 1. For Bounded Intercept/Slope and Fixed Coefficients simple returns are used, and the valuation ratios are not in logs. Further, these two restrictions only apply to valuation ratios. The initial estimation period is at least 20 years, but the OOS forecast does not start before 1926 for annual returns or December 1926 for overlapping annual returns. For the dividend-price ratio and the dividend yield, annual and overlapping annual returns are used. For the yields, quarterly returns are used.

	Model-Based Prior ΔR^2_{OS}				**CT Restrictions ΔR^2_{OS}**		
	HF	PT	LRR	HFTS	Bounded Slope	Bounded Interc. /Slope	Fixed Coeff.
Div.-Price Ratio							
Returns 1872-2011	0.01%	0.32%	1.86%		0.00%	-0.39%	1.07%
Overl. Ret. 1872-2011	1.16%	1.26%	2.86%		0.29%	2.18%	3.71%
Returns 1947-2011	7.35%	6.52%	6.93%		0.00%	4.29%	4.29%
Overl. Ret. 1947-2011	7.11%	5.94%	5.82%		0.00%	5.12%	5.12%
Dividend Yield							
Returns 1872-2011	1.02%	0.98%	2.32%		0.00%	0.93%	1.24%
Overl. Ret. 1872-2011	1.59%	1.71%	2.72%		0.00%	1.31%	1.62%
Returns 1947-2011	2.60%	5.06%	5.95%		0.00%	5.46%	5.55%
Overl. Ret. 1947-2011	2.92%	5.82%	7.38%		0.00%	6.77%	6.93%
Short Term Yield							
Ret. 1Q 1920-4Q 2011				0.31%	0.01%		
Ret. 1Q 1947-4Q 2011				2.26%	0.00%		
Long Term Yield							
Ret. 1Q 1920-4Q 2011				1.38%	-1.53%		
Ret. 1Q 1947-4Q 2011				3.33%	0.06%		
Term Spread							
Ret. 1Q 1920-4Q 2011				0.16%	0.02%		
Ret. 1Q 1947-4Q 2011				2.52%	0.00%		

Chapter 5

Conclusion

In this thesis, I investigate the link between hedge fund illiquidity and return predictability, develop a method to estimate the time-varying risk exposure and analyse the risk-shifting of hedge funds, and propose a new approach for imposing economic restrictions on parameter estimates for the purpose of predicting the equity premium.

In Chapter 2, we show that a very simple measure of hedge fund illiquidiy has in-sample and out-of-sample predictive power for returns across major asset classes. This measure is computed by averaging fund-specific return autocorrelations across a large universe of hedge funds – a high autocorrelation implies high illiquidity. We find that high hedge fund illiquidity predicts high returns and underpin this finding with a theoretical model of liquidity provision. In this model, a hedge fund acts as a market maker for a risky asset and alters its trading behaviour when faced with portfolio illiquidity constraints – an illiquid hedge fund sells and buys the risky asset for a lower price than a liquid hedge fund. In the case where the hedge fund buys the risky asset, the return reversal in the next period will be greater, and when the hedge fund sells, the return reversal will be smaller (in absolute values). These two effects combined confirm the finding from our empirical analysis: high hedge fund illiquidity forecasts high returns. The implications of this chapter are of great interest to researchers, regulators, and hedge fund investors, as so far, evidence on hedge funds having an impact on underlying asset returns has been scarce.

Information on the asset holdings of institutional investors can be highly restricted. Hence, researchers often use the variance of portfolio returns as a proxy for risk. However, hedge fund data pose a great challenge regarding the estimation of a time-varying measure of portfolio risk due to the low frequency and short time series of reported returns. In Chapter 3, I introduce a new methodology for measuring the time-varying risk of hedge fund portfolios, which is based on recent advances in the conditional variance and correlation estimation literature. The methodology provides estimates of a hedge fund's risk exposure to the core strategy and to the non-core strategies. Based on these measures, I find that the exposure to the core and non-core strategies is higher for equity long-short hedge funds which have strongly underperformed or outperformed their peers in recent months. That funds which try to catch up with their peers increase the risk of their portfolios is a well established result in the hedge fund and mutual fund tournament literature. However, the fact that outperforming hedge funds engage in risk-shifting is a new finding. The methodology for the estimation of time-varying risk of hedge fund portfolios is a contribution to the hedge fund literature with a wide applicability, and the findings about the risk-shifting of outperforming hedge funds could have interesting implications for the hedge fund industry as a whole.

In Chapter 4, the return predictability theme from Chapter 2 is revisited. As in Chapter 3, I propose a new econometric approach to address a question in empirical asset pricing. In the equity premium prediction literature, the forecasting power of standard predictors is strongly debated. Researchers have tried to improve the predictability of models with a single predictor through restrictions on the parameter estimates. I follow this path and demonstrate how to impose parameter restrictions through a well specified Bayesian framework: model-based priors are derived from consumption-based asset pricing models and imposed on single variable predictive regressions. These priors increase the OOS R^2 of up to several percentage points, which results in substantially higher expected returns for an in-

vestor with mean-variance preferences. Further, this methodology does not only improve the forecasting power of standard predictive variables, it also allows us to compare consumption-based asset pricing models based on the effectiveness of their priors.

The insights gained in this thesis lead to new research questions. In future work, I would like to explore additional channels through which hedge funds have an impact on the returns of underlying assets and whether the risk-shifting of underperforming and outperforming hedge funds can explain observed long term dynamics of investor flows in the hedge fund industry. Further, I would like to analyse, in more detail, the lessons that can be learned about consumption-based asset pricing models through the performance of their model-based priors.

Bibliography

Acharya, V., and L. H. Pedersen. 2005. Asset pricing with liquidity risk. *Journal of Financial Economics* 77:375-410.

Agarwal, V., and N. Y. Naik. 2004. Risks and Portfolio Decisions Involving Hedge Funds. *Review of Financial Studies* 17:63-98.

Agarwal, V., N. D. Daniel, and N. Y. Naik. 2011. Do Hedge Funds Manage Their Reported Returns? *Review of Financial Studies* 24:3281-3320.

Aragon, G. O. 2007. Share restrictions and asset pricing: Evidence from the hedge fund industry. *Journal of Financial Economics* 83:33-58.

Aragon, G. O., and V. Nanda. 2012. Tournament Behavior in Hedge Funds: High-water Marks, Fund Liquidation, and Managerial Stake. *Review of Financial Studies* 25:937-974.

Baker, M., and J. Wurgler. 2000. The Equity Share in New Issues and Aggregate Stock Returns. *Journal of Financial Economics* 55:2219-2257.

Bansal, R., D. Kiku, and A. Yaron. 2007. Risks for the Long Run: Estimation and Inference. Unpublished Working Paper.

Bansal, R., D. Kiku, and A. Yaron. 2010. Risks for the Long Run: Estimation and Inference. Unpublished Working Paper.

Bansal, R., and A. Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59:1481-1509.

Barberis, N., M. Huang, and T. Santos. 2001. Prospect Theory and Asset Prices. *Quarterly Journal of Economics* 116:1-53.

Basak, S., A. Pavlova, and A. Shapiro. 2007. Optimal Asset Allocation and Risk Shifting in Money Management. *Review of Financial Studies* 20:1583-1621.

Bollen, N. P. B., and V. K. Pool. 2008. Conditional Return Smoothing in the Hedge Fund Industry. *Journal of Financial and Quantitative Analysis* 43:267-298.

Bollen, N. P. B., and V. K. Pool. 2009. Do Hedge Fund Managers Misreport Returns? Evidence from the Pooled Distribution. *Journal of Finance* 64:2257-2288.

Bollen, N. P. B., and R. E. Whaley. 2009. Hedge Fund Risk Dynamics: Implications for Performance Appraisal. *Journal of Finance* 64:987-1037.

Bollerslev, T. 1987. A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return. *Review of Economics and Statistics* 69:542-547.

Bollerslev, T., and J. M. Wooldridge. 1992. Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews* 11:143-172.

Bongaerts, D., F. de Jong, and J. Driessen. 2012. An asset pricing approach to liquidity effects in corporate bond markets. Unpublished working paper.

Brown, K. C., W. V. Harlow, and L. T. Starks. 1996. Of Tournaments and Temptations: An Analysis of Managerial Incentives in the Mutual Fund Industry. *Journal of Finance* 51:85-110.

Brown, S. J., W. N. Goetzmann, and J. Park. 2001. Careers and Survival: Competition and Risk in the Hedge Fund and CTA Industry. *Journal of Finance* 56:1869-1886.

Brunnermeier, M., and L. H. Pedersen. 2009. Market Liquidity and Funding Liquidity. *Review of Financial Studies* 22:2201-2238.

Burnside, C., M. Eichenbaum, and S. Rebelo. 2011. Carry Trade and Momentum in Currency Markets. *Annual Review of Financial Economics* 3:511-535.

Campbell, J. Y. 1987. Stock returns and the term structure. *Journal of Financial Economics* 18:373-399.

Campbell, J. Y., and R. J. Shiller. 1988. The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *Review of Financial Studies* 1:195-228.

Campbell, J. Y., S. J. Grossmann, and J. Wang. 1993. Trading Volume and Serial Correlation in Stock Returns. *Quarterly Journal of Economics* 108:905-939.

Campbell, J. Y., and J. H. Cochrane. 1998. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. Center for Research in Security Prices Working Paper No. 412.

Campbell, J. Y., and J. H. Cochrane. 1999. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* 107:205-251.

Campbell, J. Y., and S. B. Thompson. 2008. Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average? *Review of Financial Studies* 21:1509-1531.

Chevalier, J., and G. Ellison. 1997. Risk Taking by Mutual Funds as a Response to Incentives. *Journal of Political Economy* 114:389-432.

Clark, T. E., and K. D. West. 2006. Using out-of-sample mean squared prediction errors to test the martingale difference hypothesis. *Journal of Econometrics* 135:155-186.

Del Negro, M., and F. Schorfheide. 2004. Priors from General Equilibrium Models for VARs. *International Economic Review* 45:643-673.

Diebold, F. X., and R. S. Mariano. 1995. Comparing Predictive Accuracy. *Journal of Business & Economic Statistics* 13:253-263.

Dimson, E. 1979. Risk measurement when shares are subject to infrequent trading. *Journal of Financial Economics* 7:197-226.

Doan, T., R. Litterman, and C. Sims. 1984. Forecasting and Conditional Projections Using Realistic Prior Distributions. *Econometric Reviews* 3:1-100.

Engle, R. F. 1982. Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica* 50:987-1007.

Engle, R. F., and K. K. Sheppard. 2001. Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH. Unpublished Working Paper.

Engle, R. F. 2002. Dynamic Conditional Correlation - a Simple Class of Multivariate GARCH Models. *Journal of Business & Economic Statistics* 20:339-350.

Engle, R. F., N. Shephard, and K. K. Sheppard. 2008. Fitting vast dimensional time-varying covariance models. Unpublished Working Paper.

Engle, R. F. 2012. Dynamic Conditional Beta. Unpublished Working Paper.

Epstein, L., and S. Zin. 1989. Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica*, 57:937-969.

Fama, E. F., and K. French. 1988. Dividend yields and expected stock returns. *Journal of Financial Economics* 22:3-25.

Fama, E. F., and K. French. 1989. Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics* 25:23-49.

Fung, W., and D. A. Hsieh. 1997. Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds. *Review of Financial Studies* 10:275-302.

Fung, W., and D. A. Hsieh. 2001. The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers. *Review of Financial Studies* 14:313-341.

Fung, W., and D. A. Hsieh. 2004. Hedge Fund Benchmarks: A Risk Based Approach. *Financial Analyst Journal* 60:65-80.

Fung, W., D. A. Hsieh, N. Y. Naik, and T. Ramadorai. 2008. Hedge Funds: Performance, Risk, and Capital Formation. *Journal of Finance* 63:1777-1803

Getmansky, M., A. W. Lo, and I. Makarov. 2004. An econometric model of serial correlation and illiquidity in hedge fund returns. *Journal of Financial Economics* 74:529-609.

Gordon, M. J. 1962. *The Investment, Financing, and Valuation of the Corporation.* Homewood, IL: Irwin.

Goyal, A., and I. Welch. 2008. A Comprehensive Look at the Empirical Performance of Equity Premium Prediction. *Review of Financial Studies* 21:1455-1508.

Gromb, D., and D. Vayanos. 2002. Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics* 66:361-407.

Gromb, D., and D. Vayanos. 2010. Limits of Arbitrage: The State of the Theory. *Annual Review of Financial Economics* 2:251-275

Hodder, J. E., and J. C. Jackwerth. 2007. Incentive Contracts and Hedge Fund Management. *Journal of Financial and Quantitative Analysis* 42:811-826.

Jylha, P., K. Rinne, and M. Suominen. 2014. Do Hedge Funds Supply or Demand Liquidity? *Review of Finance* 18:1259-1298.

Kang, N., P. Kondor, and R. Sadka. 2012. Do Hedge Funds Reduce Idiosyncratic Risk? *Journal of Financial and Quantitative Analysis* forthcoming.

Khandani, A. E., and A. W. Lo. 2011. What happened to the quants in August 2007? Evidence from factors and transactions data. *Journal of Financial Markets* 14:1-46.

Lettau, M. and S. Ludvigson. 2001. Consumption, Aggregate Wealth, and Expected Stock Returns. *Journal of Finance* 56:815-849.

Lo, A. W. 2008. *Hedge Funds: An Analytic Perspective.* Princeton, NJ: Princeton University Press.

Mamaysky, H., M. Spiegel, and H. Zhang. 2008. Estimating the Dynamics of Mutual Fund Alphas and Betas. *Review of Financial Studies* 21:233-264.

Meese, R. A., and K. Rogoff. 1983. Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample? *Journal of International Economics* 24:3-24.

Mehra, R., and E. C. Prescott. 1985. The Equity Premium: A Puzzle. *Journal of Monetary Economics* 15:145-161.

Mitchell, M., L. H. Pedersen, and T. Pulvino. 2007. Slow Moving Capital. *American Economic Review P&P* 97:215-220.

Newey, W. K., and K. D. West. 1987. A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55:703-708.

Pakel, C., N. Shephard, and K. K. Sheppard. 2011. Nuisance Parameters, Composite Likelihoods and a Panel of GARCH Models. *Statistica Sinica* 21:307-329.

Pastor, L., and R. F. Stambaugh. 2003. Liquidity Risk and Expected Stock Returns. *Journal of Political Economy* 111:642-685.

Patton, A. J., and T. Ramadorai. 2013. On the High-Frequency Dynamics of Hedge Fund Risk Exposures. *Journal of Finance* 68:597-635.

Patton, A. J., T. Ramadorai, and M. Streatfield. 2013. Change You Can Believe In? Hedge Fund Data Revisions. *Journal of Finance* forthcoming.

Pettenuzzo, D., A. Timmermann, and R. Volkanov. 2013. Forecasting Stock Returns under Economic Constraints. Brandeis Working Paper 2013/57.

Sadka, R. 2011. Liquidity risk and accounting information. *Journal of Accounting and Economics* 52:144-152.

Scholes, M., and J. Williams. Estimating betas from nonsynchronous data. *Journal of Financial Economics* 5:309-327.

Shao, J. 1989. The efficiency and consistency of approximations to the jackknife variance estimators. *Journal of the American Statistical Association* 84:114-119.

Shao, J., and C. F. J. Wu. 1989. A General Theory for Jackknife Variance Estimation. *Annals of Statistics* 17:1176-1197.

Sirri, E. R., and P. Tufano. 1998. Costly Search and Mutual Fund Flows. *Journal of Finance* 53:1589-1622.

Taylor, S. J. 1986. *Modeling Financial Time Series.* John Wiley and Sons, New York.

Thaler, R. H., and E. J. Johnson. 1990. Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice. *Management Science* 36:643-660.

Wachter, J. A. 2005. Solving Models with External Habit. NBER Working Paper No. 11559.

Wachter, J. A. 2006. A consumption-based model of the term structure of interest rates. *Journal of Financial Economics* 79:365-399.

Zellner, A. 1971. *Introduction to Bayesian Inference in Econometrics*. John Wiley and Sons, New York.