Ambiguity, Illiquidity, and Hedge Funds: An Analysis of Recent Developments and Current Research Topics in Post-Crisis Financial Markets

Dissertation
for the Faculty of Economics, Business Administration
and Information Technology of the University of Zurich

to achieve the title of
Doctor of Philosophy
in Banking & Finance

presented by
Jan Wrampelmeyer
from Bonn, Germany

approved in April 2011 at the request of
Prof. Dr. Loriano Mancini
Prof. Dr. Markus Leippold
Prof. Dr. Fabio Trojani

The Faculty of Economics, Business Administration and Information Technology of the University of Zurich hereby authorizes the printing of this Doctoral Thesis, without thereby giving any opinion on the views contained therein.

Zurich, April 6, 2011

Chairman of the Doctoral Committee: Prof. Dr. Dieter Pfaff

Table of Contents

List of Figures

Taking Ambiguity to Reality: Robust Agents Cannot Trust the Data Too Much

Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums

The Joint Dynamics of Hedge Fund Returns, Illiquidity, and Volatility

List of Tables

The Joint Dynamics of Hedge Fund Returns, Illiquidity, and Volatility

Part I

Introduction

Introduction and Summary of Research Results

1. Introduction

Starting in 2007, the most severe financial crisis since the Great Depression shook financial markets around the world. Due to the bursting of the housing bubble, banks were forced to write down hundreds of billions of US dollars in bad loans, triggering the collapse of major financial institutions and causing severe losses in international stock markets. The crisis had extensive repercussions on the real economy and the bailouts of numerous financial institutions imposed an enormous burden on national governments. This dissertation analyzes topics that have been of utmost relevance during and after the financial crisis. More precisely, the three research papers that constitute this dissertation (i) propose a robust estimation approach for economic models including uncertainty and ambiguity aversion, (ii) investigate liquidity in the foreign exchange market, and (iii) analyze the joint dynamics of hedge fund returns, illiquidity, and volatility.

Uncertainty, both within the financial sector as well as towards the financial sector, has been a critical feature of the crisis. Already the early studies of Knight (1921) and Ellsberg (1961) document that agents are not only risk averse, but also dislike ambiguity, i.e., situations in which the unknowns are unknown.[1] During the crisis uncertainty with respect to the pricing of financial securities surged when market participants realized that current valuation approaches, for example based on credit ratings, were invalid. Moreover, uncertainty about the size and location of losses aggravated the market turmoil. Consequently, uncertainty and ambiguity aversion are important aspects of economic agents' decision making, which should be taken seriously when modeling financial markets and when taking these models to the data.

A typical consequence of heightened levels of uncertainty during crisis periods are flight-to-quality effects (Caballero and Krishnamurthy, 2008) and retraction of funding liquidity. Difficulty in securing funding for business activities in turn lowers market liquidity, especially if investors are forced to liquidate positions. This induces prices to move away from fundamentals leading to increasing losses on existing positions and a further reduction of funding liquidity. Brunnermeier and Pedersen (2009) describe and model this interrelation between funding liquidity and market liquidity as a liquidity spiral. Many markets that had never before experienced any type of liquidity crisis suffered such a spiral in 2008. Therefore, contrary to investors' common perception before the crisis, evaporating liquidity is not endemic to small emerging

[1]Intuitively, risk refers to situations where the unknowns are known, while uncertainty denotes situations where the unknowns are unknown. More precisely, Knightian or model uncertainty is defined as ambiguity about the underlying data-generating process, which can derive from unknown time-varying features that agents do not understand and cannot theorize about.

economies, but liquidity risk is a common risk factor across a broad variety of markets.[2] In contrast, given that liquidity risk is invisible in normal times, academics and practitioners frequently assume that liquidity risk does not exist and hence omit this risk factor from their analyses. This is particularly true for the foreign exchange market, which is by far the world's largest financial market. However, the crisis has shown that liquidity might diminish even in the FX market leading to large losses, for instance, for carry traders who were forced to unwind positions.

During the financial crisis, uncertainty and illiquidity undermined the solvency of financial institutions culminating in the default of Lehman Brothers in September 2008. Governments around the world spent billions of dollars to support the financial industry in an effort to restore stability. Now, in the aftermath of the crisis, politicians strive for better and stricter regulation of the banking system to prevent such severe crises in the future. This unprecedented regulatory change is likely to increase the importance of the shadow banking system in general and hedge funds in particular. For instance, the implementation of the Volcker rule in the United States is likely to lead to a shift of risky trading activities from proprietary trading desks within large banks to independent hedge funds and private equity firms. While hedge funds did not play a major role during the recent financial crisis (King and Maier, 2009) this might change in the future. Therefore, it is crucial to understand how hedge fund activity interacts with market variables such as illiquidity and volatility of financial markets.

2. Summary of Research Results and Contributions

This dissertation consists of three research papers:

- Taking Ambiguity to Reality: Robust Agents Cannot Trust the Data Too Much
  (with Fabio Trojani and Christian Wiehenkamp)
- Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums
  (with Loriano Mancini and Angelo Ranaldo)
- The Joint Dynamics of Hedge Fund Returns, Illiquidity, and Volatility

Their content and contribution are summarized in the following subsections.

2.1. Taking Ambiguity to Reality: Robust Agents Cannot Trust the Data Too Much

The financial crisis has visualized the potentially enormous implications of uncertainty and ambiguity aversion for the understanding of financial markets and asset prices. Many recent papers show that ambiguity aversion helps to provide explanations for various asset pricing features and empirical puzzles arising in different areas of finance. For example, the equity premium, risk free rate, and credit spread puzzles, the large skew of implied volatility smiles of equity

[2]See, for instance, Pástor and Stambaugh (2003) as well as Korajczyk and Sadka (2008) for the equity market and Acharya, Amihud, and Bharath (2010) for corporate bonds.

index options, the predictability of bond returns, as well as stock market non-participation, underdiversification, and home-bias have been explained by investors' ambiguity aversion.[3]

Ambiguity aversion in these dynamic models is motivated by the presence of unknown time-varying features, which agents do not understand and cannot theorize about. Therefore, these models are by assumption only ideal approximations of the unknown data generating process and the potential presence of such unknown time-varying features should be considered when taking models of ambiguity aversion to real data. The existing literature largely abstracts from this important aspect and assumes that agents know the parameters of their reference model with certainty. This paper analyzes the consequences of this assumption for economic agents and model builders, who typically need to estimate a parametric model, for instance, to implement optimal robust decision rules, to quantify the equilibrium price of ambiguity or to determine bid and ask option quotes.

In these contexts, we introduce a widely applicable robust estimation approach, which is characterized by a bounded sensitivity of estimated optimal policies and general equilibrium parameters to unknown time-varying features in the underlying data generating process. We find that such a bounded influence estimation approach is key for producing (i) estimated optimal policies that are robust to unknown time-varying features and (ii) estimated equilibrium variables that are more consistent with the assumption of ambiguity aversion in general equilibrium. First, in consumption and portfolio planning problems with ambiguity aversion, small unknown time-varying features, in the conditional mean of returns or the probability structure of rare events, can generate economically relevant utility losses, which are successfully bounded by our robust estimation approach. Second, within general equilibrium economies, unknown time-varying rare event features can severely bias estimates of risk or ambiguity premiums, produced by standard estimation approaches from cross-sectional and time series information on underlying's returns and derivative prices. These biases are virtually eliminated by our bounded influence estimation approach. Third, bid-ask spreads due to ambiguity are smaller and closer to the true spreads when robustly estimating the parameters of the reference model, which reduces costly capital requirements. Finally, in a real data study on portfolio choice with ambiguous predictability, our approach uncovers predictability structures that consistently produce both (i) a larger out-of-sample utility than classical approaches and (ii) optimal portfolio weights more robust to abnormal data structures. Moreover, when focusing on real-time portfolio strategies estimated by our robust method, a mixed strategy of an agent with a realistic degree of confidence in the predictability hypothesis can produce larger utilities than the strategy implied by a random walk assumption for market returns.

Overall, these findings and the wide applicability of our robust approach to, e.g., (pseudo) maximum likelihood, generalized method of moments, and efficient method of moments settings, suggest the general usefulness of our methodology to estimate robust optimal policies and general equilibrium parameters in a broad variety of dynamic settings of ambiguity aversion. Such applications can potentially produce a number of new insights and interpretations for the growing

[3]Epstein and Schneider (2010) provide an excellent review of this literature.

literature on ambiguity aversion in finance, which largely abstracts from the implications of robust estimation in studying the consequences of ambiguity aversion for asset pricing.

2.2. Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums

Recent events during the financial crisis of 2007–2009 have highlighted the fact that liquidity is a crucial yet elusive concept in all financial markets. While there exists an extensive literature studying the concept of liquidity in equity markets, liquidity in the foreign exchange (FX) market has mostly been neglected, although the FX market is by far the world's largest financial market.

This paper fills the gap in the literature by investigating the issue of liquidity in the FX market empirically, allowing to test the theory of liquidity spirals (Brunnermeier and Pedersen, 2009) and to analyze the impact of liquidity risk on carry trade returns. To that end, we (i) accurately measure liquidity in the FX market during the crisis of 2007–2009, (ii) quantify the amount of commonality in liquidity across different exchange rates, (iii) relate FX market liquidity to measures of funding liquidity and liquidity of equity markets, and (iv) provide evidence for liquidity risk being a risk factor for carry trade returns. In addition, our analysis contributes to the growing literature on the financial crisis which tries to understand the main stylized facts to determine the causes of the recent market turmoil.

We compute benchmark FX liquidity on a daily basis using a new comprehensive data set of ultra-high frequency return and order flow data. Ranging from January 2007 to December 2009, our sample includes the financial crisis and is thus highly relevant for analyzing liquidity. By applying a variety of liquidity measures covering the dimensions of price impact, return reversal, trading cost, and price dispersion we document time series as well as cross-sectional variation in FX liquidity. We quantify the potential cost of illiquidity in the FX markets by a realistic carry trade example which shows that FX illiquidity can aggravate losses during market turmoil by as much as 25%.

Liquidity of all exchange rates decreased dramatically during the financial crisis indicating commonality in liquidity across FX rates. Such sudden shocks to market-wide liquidity have important implications for regulators as well as investors. Regulators are concerned about the stability of financial markets, whereas investors worry about the risk-return profile of their asset allocation. Decomposing liquidity into an idiosyncratic and a common component allows investors to exploit portfolio theory to reap diversification benefits with respect to liquidity risk. Therefore, we construct a time series of systematic FX liquidity representing the common component in liquidity across different exchange rates. Empirical results show that liquidity comoves strongly across currencies supporting the notion of liquidity being the sum of a common and an exchange rate specific component.

The finding of strong commonality supports the model of Brunnermeier and Pedersen (2009) which predicts comovement in liquidity of different exchange rates during liquidity spirals. To corroborate this evidence, we relate systematic FX liquidity to proxies for uncertainty as well as funding liquidity in financial markets finding that a decrease in these variables leads to lower

FX market liquidity. Moreover, market-wide FX liquidity comoves with equity liquidity which is also consistent with the presence of funding liquidity constraints during the financial crisis.

The last part of the paper investigates whether liquidity risk helps to explain daily variation in carry trade returns. Shocks to common FX liquidity are shown to be persistent and we construct a tradable liquidity risk factor by constructing a portfolio of carry trades. This novel risk factor is correlated to shocks in liquidity as well as the carry trade risk factor of Lustig, Roussanov, and Verdelhan (2010). Compared to the latter, our liquidity risk factor has a clearer and more direct interpretation following from the theory on liquidity spirals which hypothesizes that a drop in market liquidity triggers large exchange rate movements. Apart from stressing the importance of liquidity risk in the determination of FX returns, this finding supports risk-based explanations for deviations from Uncovered Interest Rate Parity (UIP) as classical tests do not include liquidity risk.

These results have several important implications. Monitoring FX liquidity on a daily basis allows central banks and regulatory authorities to assess the effectiveness of their policies. Moreover, understanding the role of liquidity and liquidity spirals helps carry traders and investors to more adequately understand the risk of their trading, which is crucial in light of the potential losses from currency crashes coinciding with liquidity spirals.

2.3. The Joint Dynamics of Hedge Fund Returns, Illiquidity, and Volatility

During the last two decades the size and importance of the hedge fund industry has increased tremendously with significant consequences for financial markets. Being part of the shadow banking system, hedge funds are largely unregulated with no lender of last resort, no discount window, no deposit insurance and no capital requirement regulation. Therefore, in light of the recent financial crisis, the role of hedge funds in financial markets is controversial and gave rise to an increasing level of scrutiny by politicians and regulators. On the one hand, hedge funds are acknowledged to be important liquidity providers which increase market efficiency and help to hold corporate management accountable. On the other hand, hedge funds have recently been blamed for high volatility as well as liquidity problems and market crashes.

This paper sheds light on the relation between hedge fund returns, volatility, and illiquidity in the equity as well as foreign exchange market. Analyzing the joint dynamics of these variables facilitates the detection of, potentially bidirectional, causalities between the market variables and hedge fund returns yielding valuable insights for hedge fund managers, funds of hedge funds, investors as well as regulators and politicians.

The empirical evidence shows that a number of hedge fund strategies respond negatively to shocks in volatility. Thus, hedge funds do not profit from higher levels of volatility in the two markets. In particular the negative relation to FX volatility is remarkable. While an exposure to FX market variables is expected for global macro or emerging market hedge funds, the results show that even hedge funds focused on equity markets, such as e.g., long/short equity funds, have significant lower returns following a shock to FX volatility. This finding has important implications for performance evaluation and compensation of managers which are largely rewarded for generating alpha. Actually, part of hedge funds profits might simply be due to FX volatility exposure which is also a valuable insight for fund of hedge funds.

In contrast, a shock to equity illiquidity has a positive impact on most hedge fund strategies, which is consistent with hedge funds being able to earn liquidity premiums during illiquid market conditions. Therefore, hedge fund managers should mainly focus on volatility exposure in their risk management process and ensure a sufficient level of funding liquidity during crisis periods to take advantage of profitable trading opportunities when they emerge.

One of the key questions in the recent political and regulatory debate is whether hedge fund activity increases volatility and illiquidity in financial markets. This paper contributes to this discussion by providing empirical evidence regarding the responses of volatility and illiquidity to shocks in hedge fund returns. While there is no clear evidence regarding illiquidity, some hedge fund styles might impact equity and FX volatility. The main positive connection between these variables is due to the trend following factors for commodities and foreign exchange with managed futures funds exhibiting the largest loading on these factors. Moreover, the algorithmic trading activities of equity market neutral funds might increase volatility. On the other hand, dedicated short bias returns tend to be followed by lower volatility corroborating the potential benefits of short selling. Thus, I find no evidence consistent with a destabilizing effect of short selling which is frequently assumed by politicians. Overall, the response of market variables to hedge fund returns is limited. However, note that the dynamics are likely to become stronger in case new regulation, e.g., more stringent capital requirements for banks, induces further growth of the shadow banking system.

Finally, the paper documents strong linkages and lead/lag relations between equity and foreign exchange volatility and illiquidity. These bidirectional spillovers indicate integration of the FX and stock market which is important for risk management. Evidence from causality tests shows that FX illiquidity and equity volatility are the main drivers in the cross-market dynamics.

References

Acharya, V. V., Y. Amihud, and S. T. Bharath, 2010, "Liquidity Risk of Corporate Bond Returns", NBER Working paper 16394, New York University, USA.

Brunnermeier, M. K., and L. H. Pedersen, 2009, "Market Liquidity and Fund Liquidity", *The Review of Financial Studies*, 22(6), 2201–2238.

Caballero, R., and A. Krishnamurthy, 2008, "Collective Risk Management in a Flight to Quality Episode", *The Journal of Finance*, 63(5), 2195–2230.

Ellsberg, D., 1961, "Risk, Ambiguity, and the Savage Axioms", *The Quarterly Journal of Economics*, 75(4), 643–669.

Epstein, L., and M. Schneider, 2010, "Ambiguity and Asset Markets", *Annual Review of Financial Economics*, 2, 315–346.

King, M., and P. Maier, 2009, "Hedge Funds and Financial Stability: Regulating Prime Brokers Will Mitigate Systemic Risks", *Journal of Financial Stability*, 5(3), 283–297.

Knight, F., 1921, *Risk, Uncertainty and Profit.* Houghton-Mifflin, Boston, USA.

Korajczyk, R. A., and R. Sadka, 2008, "Pricing the Commonality Accross Alternative Measures of Liquidity", *Journal of Financial Economics*, 87(1), 45–72.

Lustig, H. N., N. L. Roussanov, and A. Verdelhan, 2010, "Common Risk Factors in Currency Markets", Working paper, UCLA, University of Pennsylvania, and MIT.

Pástor, L., and R. F. Stambaugh, 2003, "Liquidity Risk and Expected Stock Returns", *Journal of Political Economy*, 111(3), 642–685.

Part II

Research Papers

Taking Ambiguity to Reality: Robust Agents Cannot Trust the Data Too Much

Fabio Trojani, Christian Wiehenkamp, and Jan Wrampelmeyer

This paper was presented at:

- Research Seminar at Maastricht University, The Netherlands (November 2009)
- 9^{th} Swiss Doctoral Workshop in Finance, Gerzensee, Switzerland (June 2010)
- CEPR/Study Center Gerzensee European Summer Symposium in Financial Markets, Gerzensee, Switzerland (July 2010)
- 17^{th} Annual Meeting of the German Finance Association (DGF), Hamburg, Germany (October 2010)
- Research Seminar at University of Geneva, Switzerland (November 2010)
- 60th Midwest Finance Association Annual Meeting, Chicago, USA (March 2011)

The most recent version of the paper is available at `http://ssrn.com/abstract=1668569`

Abstract
Ambiguity aversion in dynamic models is motivated by the presence of unknown time-varying features, which agents do not understand and cannot theorize about. We analyze the consequences of this assumption for economic agents and model builders, who typically need to estimate a parametric model, e.g., to implement optimal robust decision rules or to quantify the equilibrium price of ambiguity. We first show that in such contexts robust estimation methods are essential for (i) limiting the sensitivity of robust policies to abnormal time-varying features and (ii) drawing coherent inference on estimated equilibrium variables. We then introduce a general robust estimation approach, which is applicable to a wide range of economic settings with ambiguity. In the robust consumption and portfolio planning problem, a small time-varying misspecification in either the conditional mean of returns or the structure of event risk can imply substantial utility losses, which are successfully reduced by our robust estimation method. In general equilibrium economies, unknown time-varying event risk features can severely bias estimates of risk or ambiguity premiums, implied by the cross-sectional and time series information on derivative prices and underlying returns. Such biases can be virtually eliminated using our robust estimation approach. Finally, a real-data application to portfolio choice with ambiguous predictability reveals that our robust estimation method consistently yields higher out-of-sample utilities and portfolio weights that are less sensitive to particular data constellations. These findings provide new insights and interpretations for the growing literature on ambiguity aversion in finance, which largely abstracts from the link between ambiguity aversion and robust estimation.

1. Introduction

The recent financial crisis and the resulting exuberant uncertainty have dramatically shown the potentially large implications of uncertainty aversion for our understanding of financial markets and asset prices. When making decisions, investors face risk and uncertainty. While risk is the randomness that investors can understand and model, Knightian or model uncertainty denotes the ambiguity about the underlying data-generating process, which can derive from unknown time-varying features that agents do not understand and cannot theorize about.[1]

Many recent models have taken ambiguity aversion seriously, to provide potential explanations for several asset pricing features and empirical puzzles, arising in different domains of finance, including, e.g., the equity premium, risk free rate and credit spread puzzles, the large skew of implied volatility smiles of equity index options, the predictability of bond returns, as well as stock market non-participation, underdiversification, home-bias or flight-to-quality portfolio effects; Epstein and Schneider (2010) provide an excellent review of this literature. A large fraction of this research typically derives quantitative asset pricing predictions using models that are calibrated to unconditional asset prices. More recently, a few researchers have also tried to take dynamic specifications of ambiguity to the data more systematically.[2]

An important aspect, from which the literature largely abstracts, is that models of ambiguity aversion are by assumption only ideal approximations of the unknown data generating process, which in reality can contain time-varying components of unknown form that agents do not understand. Therefore, the potential presence of such unknown time-varying features should be considered to coherently take models of ambiguity aversion to real data. This paper studies the main consequences of these general features from the perspective of ambiguity averse agents and model builders who, in order to implement robust decision rules or to quantify key equilibrium parameters, have to estimate ideal models subject to unknown time-varying components.

We first show that robust estimation methods, characterized by a bounded sensitivity to such time-varying features, are essential (i) for limiting the sensitivity of estimated robust policies to abnormal features in the data generating process and (ii) for drawing coherent inference on estimated equilibrium variables. To the best of our knowledge, all classical estimation and calibration approaches applied in the literature, based on, e.g., (pseudo) maximum likelihood or generalized method of moments (GMM) methods, do not feature these basic robustness properties. Therefore, the quantitative implications derived from these methods, e.g., for optimal policies and general equilibrium parameters in economies with ambiguity, tend to highly depend on the specific form of time-varying features that might have influenced the data generating process. This is in stark contrast with the initial goal of ambiguity averse agents or model builders,

[1] At least since Knight (1921) and Ellsberg (1961), agents are known to dislike ambiguity, and various theoretical approaches to specify ambiguity aversion have been introduced, recently also for dynamic economies. See, among others, Gilboa and Schmeidler (1989), Epstein and Wang (1994), Hansen and Sargent (2001), Chen and Epstein (2002), Epstein and Schneider (2003, 2007), Klibanoff, Marinacci, and Mukerji (2005, 2009).

[2] For instance, Benigno and Nisticò (2009) estimate a VAR(1) including various macroeconomic variables and excess returns, in order to study international portfolio allocation with uncertainty. Ulrich (2009) applies pseudo maximum likelihood methods to a structural yield curve model with ambiguity aversion, while Chen, Ju, and Miao (2009) rely on maximum likelihood to estimate asset return dynamics with ambiguous predictability.

who aim to develop decision rules and policies with performance not excessively dependent on the structure of these time-varying components.

We then borrow from the vast literature on robust estimation in statistics and econometrics and introduce a general robust (bounded influence) estimation approach, implying a bounded theoretical sensitivity to time-varying features of the data generating process, which is widely applicable to a broad variety of dynamic financial models of ambiguity aversion. We apply our robust estimation approach to a number of models and quantify the resulting economic implications for estimated optimal policies and general equilibrium parameters under ambiguity aversion.

First, in consumption and portfolio planning problems with ambiguity aversion, we find that small unknown time-varying features in the conditional mean of returns or the probability structure of rare events can generate economically relevant utility losses, which are successfully bounded and reduced by our robust approach. Second, time-varying features in the context of rare events in equilibrium economies can severely bias estimates of risk or ambiguity premiums, implied by cross-sectional and time series information on derivative prices and underlying returns for classical estimation approaches. Such biases are virtually eliminated by our robust methodology. Third, in a real data application to portfolio choice with ambiguous predictability, our approach uncovers predictability structures that consistently produce (i) a larger out-of-sample utility and (ii) optimal portfolios that are less dependent on abnormal data structures. In addition, a real-time mixed strategy, estimated by our robust method for investors with a realistic degree of confidence in the predictability hypothesis, produces a higher utility than a simple random walk assumption for market returns.

Overall, our findings indicate that the economic costs of being indifferent to robustness at the model estimation stage can be large, when ambiguity is present: Our examples of concrete ambiguous economies show that they can even dominate the potential costs of non-robust optimal policies in contexts where the benchmark model is common knowledge. In these settings, our robust (bounded influence) approach successfully limits the potential damaging impacts of unknown time-varying features in the data generating process, while still ensuring a quite efficient performance under ideal (unambiguous) model conditions.

More generally, our robust estimation approach can potentially provide a new way to interpret and study the findings of the growing literature on ambiguity aversion in finance, which typically abstracts from a concern for robustness at the model calibration and estimation stage. In this literature, an impressive number of relevant asset pricing question has been addressed, using the common ambiguity aversion denominator. Maenhout (2004) and Trojani and Vanini (2002), among others, find that ambiguity aversion can increase the equity premium and lower the risk free rate, for reasonable choices of risk aversion, while Leippold and Trojani (2008) show that the interaction of learning and ambiguity additionally explains the excess volatility of stock returns. Liu, Pan, and Wang (2005) demonstrate that uncertainty about rare events plays an important role for the smirk patterns of option markets and Boyle, Feng, Tian, and Wang (2008) select robust pricing kernels in incomplete markets to robustly price new derivatives. Trojani and Sbuelz (2008) highlight that time-varying ambiguity aversion can be linked to countercyclical stock returns and volatilities. Cao, Wang, and Zhang (2005) and Trojani and Vanini (2004) show that dispersion in uncertainty aversion across investors can produce an endogenous stock market participation. Routledge and Zin (2009) explain that ambiguity can limit market makers hedging ability and thus reduces liquidity, while Easley and O'Hara (2009) argue that non-participation can be linked to the ambiguity aversion of rational traders, explaining the absence of trading during the 2007–2009 financial crisis. Epstein and Miao (2003) and Uppal and Wang (2003) demonstrate that different degrees of ambiguity aversion across stocks or economies can imply home bias and underdiversification effects. Gagliardini, Porchia, and Trojani (2009) show that ambiguity aversion can explain deviations from the expectations hypothesis in bond markets, while Chen, Ju, and Miao (2009) conclude that ignoring the uncertainty of the link between stock returns and dividend-price ratios can lead to sizable welfare costs in dynamic asset allocation. Caballero and Krishnamurthy (2008) find that a larger ambiguity during financial crises can imply flight-to-quality effects. Boyarchenko (2010) shows that ambiguity about the quality of accounting signals and asset dynamics helps explain credit default swap spread dynamics. Drechsler (2009) demonstrates that uncertainty about economic fundamentals plays a central role in explaining equity premiums, return volatilities and option skews. Benigno and Nisticò (2009) show that part of the US home bias can derive from a hedging demand against long-run real exchange rate ambiguity. Ju and Miao (2010) generate procyclical (countercyclical) variation of dividend-price ratios (equity premiums) in a model with ambiguous consumption and dividends. Miao and Wang (2010) consider uncertainty in relation to option exercise, when modeling real investment and exit problems. Finally, Ulrich (2009) finds that uncertainty aversion significantly affects equilibrium inflation premiums.

Our work is also related to the large robustness literature in econometrics and statistics, showing that many broadly used (classical) estimators tend to be highly sensitive to particular time-varying features in the data generating process. Theoretical and empirical evidence on this failing robustness is provided, among others, by Künsch (1984) for autoregressive models, Bustos and Yohai (1986) for ARMA models, Krishnakumar and Ronchetti (1997) for simultaneous equation models, Muler and Yohai (1999) for ARCH models, Ronchetti and Trojani (2001) and Gagliardini, Trojani, and Urga (2005) within Generalized Method of Moments models, Genton and Ronchetti (2003) for Indirect Inference methods, Ortelli and Trojani (2005) for Efficient Method of Moments settings, Sakata and White (1998) and Mancini, Ronchetti, and Trojani (2005) for GARCH type volatility models, Mancini and Trojani (2010) for tail estimation and La Vecchia and Trojani (2010) for diffusion models; see also Hampel, Ronchetti, Rousseeuw, and Stahel (1986) for an overview on the bounded influence robust approach in statistics.

The remainder of the paper is organized as follows. Section 2 introduces a simplified example of ambiguity aversion in conjunction with parameter estimation, which clarifies the main issues and objectives of our analysis. Section 3 introduces our general estimation approach and presents a class of robust (bounded influence) estimators applicable to a wide variety of dynamic models with ambiguity aversion. Section 4 quantifies the implications of robust estimation for a number of relevant applications of ambiguity aversion, including the portfolio choice problem with rare events, the estimation of equilibrium risk and ambiguity premiums for jumps, and the pricing of options under ambiguity. Section 5 produces a thorough empirical study on dynamic portfolio choice with ambiguous predictability and quantifies the added economic value of our robust estimation approach in this context. Section 6 concludes.

2. Ambiguity Aversion and Robust Estimation Model Setting

A large number of recent approaches incorporating ambiguity aversion in dynamic financial decision making specifies investors' preferences using different types of max-min expected utility criteria, which imply optimal policies robust to the potential misspecification of a given benchmark model. In these settings, agents explicitly recognize that (i) models are only approximations to the unknown data generating process (DGP) and that (ii) data can be generated from a set of unspecified alternatives to the benchmark model, which are difficult to detect statistically. Epstein and Schneider (2007, 2008), for instance, emphasize that the true DGP can contain small time-varying, unknown features, which the agent does not understand and cannot even theorize about: The time-varying nature and the small size of these components can make their statistical identification very hard, if not impossible, given realistic amounts of data. The literature typically assumes that the parameters of the benchmark reference model available to ambiguity averse agents are known, so that they do not need to be estimated. Therefore, the implications of robust estimation for the utility of the policy of ambiguity averse agents are largely unstudied. Similarly, the asset pricing predictions of ambiguity aversion in general equilibrium economies are typically derived from real data calibrations based on non-robust estimation methods. Thus, the potential consequences of robust estimation for quantitative predictions of ambiguity aversion in general equilibrium are also largely unexplored. How large can these effects quantitatively be and which estimators should be used in a context of ambiguity aversions, in order to estimate models where the true DGP can contain small time-varying, unknown components? We introduce and motivate these important questions in the context of a simplified robust portfolio choice setting, based on a geometric Brownian motion (GBM) as benchmark for the market return dynamics. In later sections, we extend the main insights of this example to a general robust estimation approach that can handle a broad variety of models proposed in the literature.

2.1. Robust Portfolio Choice in a Simplified Model

Market returns follow a process that is a small perturbation of a GBM reference model. For simplicity, we assume that the perturbation affects only the drift of returns by a small, potentially time-varying, component u_t:

$$\frac{dS_t}{S_t} = \mu dt + \sigma(dB_t + u_t dt), \tag{1}$$

where $\mu \in \mathbb{R}$, $\sigma > 0$ and B is a standard Brownian motion. This drift distorted approach has been proposed by Anderson, Hansen, and Sargent (2000) and is a convenient way to model misspecification of the conditional mean of returns. Extensions that model a potential misspecification of other features of the conditional return distribution, like, e.g., conditional skewness or kurtosis, are studied later in this paper.

The specific form of the time-varying component u_t is unknown to investors, who only know that u_t is small in a statistical sense, meaning that the transition densities of GBM and the perturbed process are similar and difficult do distinguish statistically. Perturbation (1) implies the following distorted wealth dynamics relative to the benchmark GBM case:

$$dW_t = [W_t(r + \pi_t(\mu - r)) - C_t]\, dt + \pi_t \sigma W_t (dB_t + u_t dt), \tag{2}$$

where r denotes the riskless interest rate, π_t the fraction of wealth invested in the market index and C_t agent's consumption. Borrowing from insights in Anderson, Hansen, and Sargent (2000), Maenhout (2004) accounts for unspecified drift adjustments of the above type and extends Merton's (1969) portfolio planning problem by a penalization reflecting a desire for robustness. The ambiguity averse agent follows a worst case approach and selects policies that tend to avoid negative utility surprises due to model misspecification. The value function of the robust optimization problem reads:

$$V(W_0) = \sup_{\{C_t, \pi_t\}} \inf_{\{a\}} E^a \left[\int_0^T \left(U(C_t) + \frac{1}{2\Psi} (\pi_t \sigma W_t a)^2 \right) dt \right], \tag{3}$$

subject to (2), where the agent's time preference β is incorporated in the utility function $U(C_t)$. Despite the concern for robustness, the assumed benchmark model is still the best possible representation of the data. This feature is reflected by the penalty term in the objective function (Equation (3)), which constrains the agent to focus on alternative DGPs not too far away for the assumed model, when seeking to minimize the utility impact of a misspecification. The possibly state dependent parameter Ψ controls the amount of robustness the agent seeks: $\Psi \to 0$ corresponds to the classical Merton problem, as it enforces $a = 0$ as the only viable choice. On the contrary, the larger Ψ the more the agent is concerned with robustness and is willing to consider a larger set of alternatives in her worst case utility optimization. Assuming constant relative risk aversion (CRRA) utility, Maenhout (2004) shows that choosing $\Psi(W,t) = \frac{\vartheta}{(1-\gamma)V(W,t)}$ allows to analytically solve for the robust value function and optimal policies, which are given by:

$$C_t^* = \frac{\chi}{e^{-\chi \cdot (T-t)}} W_t \quad \text{and} \quad \pi_t^* = \frac{1}{\gamma + \vartheta} \frac{\mu - r}{\sigma^2}, \tag{4}$$

where γ measures relative risk aversion, ϑ is the amount of robustness the agent seeks and $\chi = \frac{1}{\gamma} \left(\beta - r(1-\gamma) - \frac{1-\gamma}{2(\gamma+\vartheta)} \left(\frac{\mu - r}{\sigma} \right)^2 \right)$.

In order to implement the optimal policies, the robust agent needs to know parameters μ, σ of the benchmark GBM model for market returns. The literature on ambiguity aversion typically assumes these parameters to be known and does not quantify the potential impact of different estimation approaches on the effective utility implied by the optimal policies of a robust agent. In reality, however, model parameters have to be estimated and a suitable estimation procedure has to be chosen to estimate a model with a sufficient degree of precision. In presence of ambiguity and small time-varying unknown features in the true DGP this can be a complicated task, e.g., because maximum likelihood type estimators are no longer optimal and can even imply biased or very inefficient estimators, depending on the specific form of perturbation u_t in Model (1). These non-robust features of maximum likelihood type estimators are well-known in the statistical and econometric literature and have led several authors to consider robust statistical theories, in which different forms of robust maximum likelihood type estimators have been proposed.[3] In a context where data have been generated by a DGP that is only approximately described by the assumed model, such robust estimators are natural potential alternatives to maximum likelihood, since they can ensure a more consistent estimation of a benchmark model in presence of even small (time-varying) unknown features of the true DGP.

2.2. The Added Value of Being Ambiguity Averse When the Benchmark Model is Known

Assume that the parameters μ, σ of the GBM part in Equation (1) are known. What is the added value for an economic agent to be ambiguity averse, as opposed, e.g., to a pure expected utility optimizer? Intuitively, the expected utility agent will produce optimal policies with the highest utility if returns exactly follow a GBM. At the same time, she will produce policies with a lower utility if the true DGP contains some unknown time-varying features as in Equation (1): The utility loss relative to the GBM case will depend on the specific form of distortion u_t. In contrast, the ambiguity averse agent will consider ex-ante optimal policies that account for the potentially negative impact of small time-varying components of unknown form in the return dynamics. Therefore, she will select robust policies with a lower utility if returns follow indeed a GBM, but which at the same time imply a performance not excessively dependent on the specific form of u_t, if the DGP is indeed different from a GBM. The ambiguity averse agent pays an ambiguity premium for following a suboptimal policy in the case where returns exactly follow the given benchmark model, but she also ensures a satisfactory degree of utility for the case where the DGP indeed contains small time-varying components of unknown form. It follows that the utility of the robust policy is not excessively dependent on which particular distortion in Equation (1) can be produced by malevolent nature. In contrast, the expected utility maximizer makes the success of her decision rule mainly dependent on luck. The utility of her optimal policy is largely exposed to the specific form of potential distortion u_t selected by nature to generate the data, which is a feature completely beyond the agent's control.

[3]Starting with the seminal work of Huber (1964) and Hampel (1974), robust statistical theories have been extensively developed in the literature; see e.g. Huber (1981) for a review. While these classical papers consider the iid context, more recently, robust estimators have been proposed for a variety of estimation problems in a general dynamic time series setting. For instance, Sakata and White (1998) robustly estimate GARCH type models, Ronchetti and Trojani (2001) introduce robust GMM, Genton and Ronchetti (2003) robustify indirect inference, Ortelli and Trojani (2005) develop robust efficient methods of moments (EMM) estimators, while robust estimators for diffusions are studied by La Vecchia and Trojani (2010).

2.3. The Added Value of Being Ambiguity Averse in a Concrete Example

It is useful to quantify the relative performance of the policies of expected utility and ambiguity averse agents in realistic applications. To this end, we consider a simple specification for the unknown time-varying component u_t in Equation (1).

2.3.1. Modeling the Drift Distortion

Without loss of generality, we assume that the majority of the data is generated by a GBM with $\mu = 0.06$ and $\sigma = 0.15$. At the same time, we allow the conditional expected market return of a small fraction of the daily observations to follow a time-varying drift process, which very infrequently distorts the constant benchmark expected return μ over a time period of $H = 6$ years. Formally, the distorted drift is specified as:

$$\mathbb{E}\left(\frac{dS_t}{S_t}\right) := \mu + \sigma u_t = \mu + \eta_t(\alpha_t - \mu) \quad ; \quad t \in [0, H], \tag{5}$$

where η_t is an iid Bernoulli process independent of Brownian motion B, such that $P(\eta_t = 1) = 0.03$, and α_t is a piecewise constant function of time defined by:

$$\alpha_t = \begin{cases} -0.125 & t \in [0, 0.21H) \\ 0.09 & t \in [0.21H, 0.54H) \\ -0.0025 \cdot 252 & t \in [0.54H, 0.58H) \\ 0.03 \cdot 252 & t \in [0.58H, 0.75H) \\ -0.0075 \cdot 252 & t \in [0.75H, 0.83H) \\ -0.025 \cdot 252 & t \in [0.83H, H]. \end{cases} \tag{6}$$

Therefore, in about 3% of the cases, the constant expected return in the geometric Brownian motion setting can be shifted upwards or downwards in a time-varying way. For instance, in the sixth year, it is possible in about 3% of the days that the daily expected market return is -2.5%, which can model, e.g. rare negative news on financial markets that can temporarily strongly affect expected returns. The unconditional mean of the drift distortion in Equation (5) is zero, i.e., the distortion is symmetric: it alters the tail behavior of market returns but not their unconditional mean. For illustration purposes, the top left panel of Figure 1 depicts the return density of Model (1) for the drift distortion in Equation (5), together with the density of a GBM with $\mu = 0.06, \sigma = 0.15$, suggesting a small discrepancy between these two unconditional distributions. Moreover, sample realizations for the distorted process (top right panel) and the decomposition into clean data (bottom right panel) and the time-varying drift component (bottom left panel) visualize the difficulty to distinguish the two processes.

[Figure 1 about here.]

Even if the tail behavior of the two distributions is different, without good prior knowledge, it is very difficult to determine the specific structure of the time-varying component from the data.

This is due to the fact that the perturbation in Equation (5) implies infrequent deviations from the benchmark model and is dominated by a substantial time-varying part. Thus, investors cannot fully understand or even theorize about these complicated time-varying features, but they suspect a potential misspecification of unknown form in their benchmark model, which motivates their use of robust decision rules.

2.3.2. The Added Value of Ambiguity Aversion Given a Known Benchmark Model

Investors observe returns for the first five years, and collect a sample of daily returns $\Delta \log S_t$, for $t = 1, \ldots, 252 \cdot T$, where $T = 5$. At the end of the fifth year, they select optimal consumption and portfolio policies according to Equation (4), using an investment horizon of one year. We first assume that parameters $\mu = 0.06$ and $\sigma = 0.15$ are common knowledge, so that investors do not have to care about their estimation. This is the situation typically studied in the literature. Based on this parameter choice, we first determine the optimal policy for (i) the expected utility agent and (ii) the ambiguity averse agent, in dependence of robustness parameter ϑ. We then compute the certainty equivalent wealth implied by these ex-ante policies in the final year for two different possible scenarios: First, in a scenario where the data follow a GBM and, second, in a scenario where the GBM is distorted according to Model (5). Of course, these are just two possible ex-post realizations of time-varying component u_t, and a whole variety of small time-varying drift misspecifications u_t could be considered. However, the main message of our analysis, i.e., the excessive sensitivity of the performance of non-robust policies to the specific shape of u_t, would not change. Therefore, we focus for brevity on drift distortions related to the misspecification in Equation (6). Panel (a) of Table 1 summarizes the first set of results.

[Table 1 about here.]

Columns (1) and (2) compute the ex-ante certainty equivalent wealth for the different agents.[4] The expected utility agent has an ex-ante wealth of 1,047, which is higher than the ex-ante wealth of ambiguity averse agents, which ranges between 690 and 919, depending on ϑ. The lower ex-ante wealth of the ambiguity averse agent reflects her ex-ante worst-case approach in accounting for potential misspecification, which is stronger for larger values of ϑ. Columns (3) and (5) compute the ex-post certainty equivalent wealth of the same agents in the case where in the final year the DGP is indeed a GBM. Since ex-ante and ex-post the DGP is the same and follows a GBM, the expected utility agent achieves ex-post the same highest equivalent wealth as she was expecting. Interestingly, despite the worst-case ex-ante approach in fixing the optimal policies, the ex-post scenario-dependent equivalent wealth of all robust agents is only marginally lower than for the expected utility maximizer. It follows that for the selected drift-distortion scenario the ambiguity premium paid by ambiguity averse agents to apply robust decision rules is small. Column (4) and (6) reproduce the equivalent wealth in a scenario where ex-post the DGP contains a contaminated drift (Equation (5)). The expected utility maximizer has an ex-post

[4] In Panel (a), Columns (1) and (2), (3) and (5) as well as (4) and (6) are identical because parameters μ and σ are not estimated.

equivalent wealth of 926, which is about 12% lower than in the case with no contamination. At the same time, the equivalent wealth of the robust agents ranges from 934 to 988 depending on the ex-ante degree of robustness ϑ. Relative to the case with no contamination, the reduction in equivalent wealth is lower than for the expected utility maximizer and it is only about 6% for the more ambiguity averse agent $\vartheta = 3$. Overall, these results quantify the potential added value of a robust approach to portfolio choice in the above setting: While in absence of misspecification a robust decision rule produces a small loss in utility, in presence of a small time-varying distortion it can avoid a good fraction of the large losses of about one fourth of certainty equivalent wealth suffered by the expected utility agent.

2.4. Determination of the Benchmark Model in Presence of Ambiguity: (Pseudo) Maximum Likelihood

2.4.1. (Pseudo) Maximum Likelihood Estimation

Ideally, the robust agent would like the estimated reference model to approximate as good as possible, given the available data, the distribution implied by the unknown DGP, because this would make the robust decision rule applied to the estimated reference model more effective. A natural starting point to achieve this goal is the standard pseudo maximum likelihood (PML) estimator $(\hat{\mu}, \hat{\sigma})$ implied by the GBM assumption for the reference model in Equation (1), which is given by the solution of the equations:

$$\sum_{t=1}^{T} \xi\left(\frac{\Delta \ln S_t - \hat{\mu}}{\hat{\sigma}}\right) = 0, \tag{7}$$

$$\sum_{t=1}^{T} \chi\left(\frac{\Delta \ln S_t - \hat{\mu}}{\hat{\sigma}}\right) = 0, \tag{8}$$

where $\xi(z) = z$ and $\chi(z) = \xi(z)^2 - 1$. PML estimators are the most efficient unbiased estimators in the case where the benchmark GBM model is correctly specified.[5] Moreover, if the benchmark model is misspecified in a time-invariant way, i.e., when in-sample and out-of-sample DGP are identical, PML estimators have the convenient property of minimizing the statistical discrepancy between the unknown stationary DGP and the reference model, where discrepancy is measured by Kullback-Leibler (KL) divergence.[6] Thus, in these settings PML estimators are very convenient estimators from the perspective of an ambiguity averse agent. However, in settings of uncertainty where the true DGP can contain small time-varying, unknown features, leading to different in-sample and out-of-sample DGPs, PML estimators lose their theoretical optimality and can imply quite dramatic losses in accuracy. Intuitively, this feature arises because of the failing robustness of these estimators, which generates a large sensitivity of many PML point estimates to even small fractions of particular observations in a random sample.

[5]In this case, maximum likelihood and PML estimators are identical, of course.

[6]In such a setting, the agent could also hope to eventually learn at least some features of the exact type of distortion present in the data.

2.4.2. Ambiguity Aversion and PML Estimation of the Reference Model

As in the previous section, we can quantify in more detail the relation between PML estimation and ambiguity aversion in the context of robust portfolio choice, for the case where the DGP for market returns follows the dynamics specified in Equation (1) with a potential drift distortion (Equation (5)). Panel (b) of Table 1 summarizes the results. Column (3) shows that the ex-post loss in utility due to model estimation using optimal maximum likelihood estimators is small: When returns follow a GBM in-sample and out-of-sample, the loss in equivalent wealth is only about 6% for the expected utility agent and between 1% and 5% for the ambiguity averse agent. These losses are much smaller than the losses due to pure misspecification in Columns (4) and (6) of Panel (a), where estimation risk is completely absent. These findings indicate that (i) pure estimation risk due to PML estimation is economically small in the given setting, when data perfectly follow the benchmark model in- and out-of-sample, and that (ii) robust decision rules tend to control implicitly also some degree of risk due to model estimation: They produce smaller utility losses relative to the expected utility agent and produce overall a larger certainty equivalent wealth.

Column (5) shows that larger utility losses arise when in-sample data follow a distorted GBM. This is a natural finding, because estimator (7)-(8) is not the most efficient (maximum likelihood) estimator under a distorted DGP, and therefore leads to less precise parameter estimates. Compared to a setting where the in-sample DGP is GBM, the loss in equivalent wealth of the expected utility agent almost doubles to about 10%, and the one of the ambiguity averse investor is between 1% and 7.5%. Apparently, robust policies still tend to control fairly well also the estimation risk of PML estimators applied to a distorted in-sample DGP: They again imply lower utility losses and a larger certainty equivalent wealth overall, relative to the policies of the expected utility optimizer. Columns (4) and (6) show that the largest losses in utility arise when in-sample and out-of-sample DGP are different, because of the presence of unknown small time-varying features in the data, generated by drift distortion (5). Column (4) considers the case where the in-sample DGP is GBM, so that the optimal estimator is maximum likelihood estimator in Equations (7)-(8). In this case, the loss in certainty equivalent wealth of the expected utility agent is about 35%. Robustification of the optimal policy can reduce this loss up to 87% in the case $\vartheta = 3$. These results indicate that robust decision rules are even more crucial when the reference model is estimated, also under ideal, i.e., maximum likelihood conditions: Compared to the results in Column (4) of Panel (a), the increase in ex-post certainty equivalent wealth of a robust policy for $\vartheta = 3$, relative to the expected utility policy, is about 7% in Column (4) of Panel (a), while it amounts to about 57% in Column (4) of Panel (b). Column (6) considers the case where also the in-sample DGP is different from a GBM. Intuitively, this is the setting where the largest losses in certainty equivalent wealth arise, because (i) the PML estimator is less efficient than under ideal GBM conditions and (ii) in-sample and out-of-sample DGP can be affected by different unknown time-varying features, as specified by Equation (5). The loss in certainty equivalent wealth of the expected utility agent is about 48%, and the one of the ambiguity averse agent ranges between 12% and 46%! Compared to Column (4), in Column (6) a good fraction of certainty equivalent wealth is lost exclusively because PML estimator (7)-(8) is applied to a distorted in-sample DGP: The utility loss of an expected utility maximizer in Column (6) relative to Column (4) is about 21%, while the one of an ambiguity averse agent ranges between 7% and 20%.

Overall, this evidence shows that robust decision rules can only moderately control the utility impact of unknown small time-varying features in the data, when models are estimated by non-robust PML estimators of the form in Equation (7)-(8): The percentage utility loss caused by non-robust PML estimation in presence of small time-varying features (the differences between Columns (6) and (4) of Panel (b)) is at least of the same order of magnitude as the loss generated when the reference model is known (the differences between Columns (4) and (3) of Panel (a)). These results indicate a large sensitivity of the performance of robust optimal policies estimated by PML estimators to small, unknown time-varying features in the data. The recent literature on ambiguity aversion in finance largely disregards this aspect, which is a main focus of this paper.

2.5. Determination of the Benchmark Model in Presence of Ambiguity: Robust (Pseudo) Maximum Likelihood

How can an ambiguity averse agent incorporate a concern for small time-varying DGP components of unknown form at the stage of the estimation of a benchmark model? Intuitively, since some features of the data are likely slightly different for in-sample and out-of-sample DGPs, it is a wise robust strategy to consider estimators that are not too sensitive to particular features of a limited fraction of observations in a random sample. Many pseudo maximum likelihood type estimators are well-known to be extremely sensitive to a few observation in sample. Therefore, a number of authors in the statistics and econometrics literature has considered robust statistical theories, in which different robust estimators have been proposed. In a context where data have been generated by a DGP that is only approximately described by a model, such estimators are natural potential alternatives to pseudo maximum likelihood estimators, because they can ensure a more consistent estimation of a benchmark model, in presence of small (time-varying) unknown features of the DGP.

2.5.1. Robust Estimation: Huber's Proposal

Let $\{\Delta \log S_t : t = 1, \ldots, 252 \cdot T\}$ be the sample of daily observations from distorted DGP (1)-(5) and denote by $f(x)$ the unknown density of standardized returns $(\Delta \log S_t - \mu)/\sigma$. The reference model density for standardized returns under a GBM process is the standard normal density $\phi(x)$. We denote by $F(x)$ and $\Phi(x)$ the corresponding distributions, where F is by assumption in a small neighborhood $\mathcal{U}$ of Φ. Huber (1964) embeds ML estimators into a broader class of M–estimators, which are defined as implicit solutions of the following equations:

$$\sum_{t=1}^{T} \psi\left(\frac{\Delta \log S_t - \mu}{\sigma}\right) = 0, \tag{9}$$

for some function $\psi(x) \in \mathbb{R}^2$. Within this class, he motivates a number of robust estimating functions ψ, which solve (worst-case) minimax problems of the form:

$$\inf_{\psi} \sup_{F \in \mathcal{U}} \mathrm{Var}(\psi, F), \tag{10}$$

where $\mathrm{Var}(\psi, F)$ is the asymptotic variance, under distribution F, of the M-estimator associated with estimating function ψ. The robust estimator implied by an estimating functions solving Problem (10) explicitly incorporates a concern for ambiguity at the estimation stage, by ensuring that the largest asymptotic variance of an estimator over neighborhood $\mathcal{U}$ around benchmark model Φ is minimized. In the context of our model setting in Equations (1) and (5), this worst case estimator is the well-known Huber estimator $(\tilde{\mu}_R, \tilde{\sigma}_R)$, defined implicitly as the solution of:

$$\sum_{t=1}^{T} \xi_c \left(\frac{\Delta \ln S_t - \tilde{\mu}_R}{\tilde{\sigma}_R} \right) = 0, \tag{11}$$

$$\sum_{t=1}^{T} \chi_c \left(\frac{\Delta \ln S_t - \tilde{\mu}_R}{\tilde{\sigma}_R} \right) = 0, \tag{12}$$

where $\xi(z) = \min(c, \max(-c, z))$ and $\chi(z) = \xi^2(z) - E_\Phi \xi^2$, for some constant c determining the tradeoff between robustness and efficiency.[7]

2.5.2. Ambiguity Aversion and Robust PML Estimation of the Reference Model

We can quantify in more detail the added value of robust estimation in the context of the robust portfolio choice problem in Equation (3), for the case where the DGP for market returns following the distorted dynamics in Equations (1) and (5). Panel (c) of Table 1 summarizes the results. Column (3) and (4) of Panel (c) show that when return data used for estimation follow a GBM, the ex-post equivalent wealth of both expected utility and robust agents is only slightly lower than the equivalent wealth in Columns (3) and (4) of Panel (b). This finding implies that using robust estimators instead of ML estimators, when market returns exactly follow the ideal GBM model, produces only a small loss in efficiency. The situation is very different in Column (5) and especially in Column (6), which represents the realistic case where market return data used for estimation follow a distorted DGP: When comparing Column (6) of Panels (c) and (b), we see that the increase in certainty equivalent wealth is about 13% for the expected utility agent, while it ranges between 4% and 12% for the ambiguity averse agents. The same comparison also shows that an agent using classical PML estimates must consider a wider set of alternatives (controlled by the choice of ϑ) in order to enjoy the same level of ex-post utility as someone that uses robust estimates. Initially, this cannot be in the agent's interest, though, given that

[7]The underlying idea of Equation (12) is that the parameters are not estimated by least squares, but by minimizing some non-linear function of the residuals subject to the constraint in Equation (12). The latter states that the estimate of the uncontaminated residual variance, $\frac{\xi^2(z)}{T-1}$, is equal to the true variance of the theoretical stochastic structure, after accounting for potential outliers. In applications, c is usually chosen such that the robust estimator achieves 95% efficiency if the data have actually been generated by the reference model. We follow this convention throughout this paper.

ex-ante expected utilities are negatively related to the agent's degree of ambiguity aversion.

Overall, these findings show that both robust decision making and robust estimation are necessary ingredients to ensure that the ex-post utility of an optimal strategy is largely independent of potential small model distortions, selected by nature to generate the data. Since agents cannot know the exact structure of the transition density of market returns, they should choose a modeling and estimation strategy that performs well also in realistic least favorable circumstances. If they do not do so, they make the success of their policy largely dependent on luck, i.e., the specific form of DGP selected by nature, which is a feature completely beyond the agent's control.

2.6. The Added Value of Robustness When the Perturbation is Reversed

In the previous sections, the process for drift distortion u_t implied by the chosen specification in Equation (6) is such that out-of-sample average returns tend to be lower than in-sample average returns, thus creating a portfolio choice setting in which in-sample PML estimates tend to imply an overly optimistic estimation of future investment opportunities. It is important to realize that the excessive sensitivity to model perturbations of optimal policies estimated by PML is not generated by the chosen form of drift distortion in Equation (6), but it is rather an intrinsic feature of policies estimated by non-robust PML methods in the above context. In order to illustrate this important point, we perform the above analysis for a reversed drift distortion process $\alpha_t' = -\alpha_t$, so that in-sample average returns are lower than out-of-sample average returns, leading to more favourable investment opportunities out-of-sample:

$$\mathbb{E}\left(\frac{dS_t}{S_t}\right) := \mu + \sigma u_t = \mu + \eta_t(\alpha_t' - \mu) \quad ; \quad t \in [0, H], \tag{13}$$

where, as before, η_t is an iid 0-1 process independent of Brownian motion B, such that $P(\eta_t = 1) = 0.03$. Results are summarized in Table 2.

[Table 2 about here.]

When the reference model parameters are common knowledge, Panel (a) of Table 2 shows that, in contrast to the previous example and as suggested by intuition, the expected utility maximizer is slightly better off than the robust agent, when the out-of-sample DGP is contaminated, with a largest improvement in certainty equivalent wealth of about 5% with respect to the most ambiguity averse agent. Therefore, in absence of parameter estimation this lucky scenario selected by nature tends to favor expected utility relative to ambiguity averse agents. However, the situation is reversed when reference model parameters have to be estimated. For instance, maximum likelihood estimates in Panel (b) always imply a lower utility of expected utility agents relative to ambiguity averse agents. Moreover, when both in- and out-of-sample data are distorted, the robust estimates in Panel (c) clearly improve on the performance of policies implied by PML estimators: For instance, an expected utility agent improves his certainty equivalent wealth by about 12%, simply by using robust instead of non-robust estimators.

Overall, we again find that robust estimators and robust policy rules are both necessary, in order to produce out-of-sample utilities not excessively dependent on the form of unknown time-varying drift distortion potentially contaminating the given benchmark model.

3. A Generally Feasible Solution to the Robust Estimation Problem

As the previous examples have shown, PML estimators applied to a DGP featuring unknown time-varying features can imply out-of-sample utilities of estimated robust policies that are heavily dependent on the specific form of the model misspecification selected by nature. This feature is in evident contrast with the initial goal of ambiguity averse agents, who aim to develop robust optimal decision rules with performance not excessively dependent on the structure of such unknown time-varying features. In the context of the simple GBM example of Section 2, Huber's (1964) max-min robust estimator was shown to produce better results than PML estimators, leading to more robust and stable utilities of estimated optimal policies across different forms of potential small perturbations of the underlying DGP. In this section, we extend Huber's (1964) approach to a broader class of robust estimators, which can be used to study more general settings and models of ambiguity aversion, including, e.g., non-Gaussian returns, predictability and stochastic volatility features.

3.1. Bounded Utility Sensitivity by Bounded Estimator Sensitivity

Consider an ambiguity averse agent with reference model DGP P_{θ_0}, $\theta_0 \in \Theta$, and a fixed degree of ambiguity aversion. The optimal policy and the worst case utility implied by the agent's max-min optimal choice depend on the reference model parameter and are denoted by $C(\theta_0)$ and $V(\theta_0)$, respectively. This utility defines a lower bound for the actual utility of the robust optimal policy with respect to a given family $\mathcal{U}$ of potential distortions $u := \{u_t\} \in \mathcal{U}$ of reference model P_{θ_0}, which are taken into account by the ambiguity averse agent in solving her max-min optimization problem. We assume that $C(\cdot)$ and $V(\cdot)$ are smooth (differentiable) functions of reference model parameter θ. This assumption is typically satisfied by most (dynamic) models of ambiguity aversion.

Assume now that in-sample data are generated by a reference model θ_0 distorted by a candidate model contamination $u \in \mathcal{U}$, implying an in-sample DGP $P_{\theta_0,u}$. A necessary robustness requirement is that estimated robust optimal policies and worst case utilities have bounded sensitivity with respect to the form of potential distortion u: An unbounded sensitivity would imply optimal decision rules with a utility performance that can be arbitrarily influenced by even a small such model distortion. Let $\tilde{\theta}(P_{\theta_0,u})$ be the asymptotic value of an estimator applied to distorted data $P_{\theta_0,u}$. Using a simple plug-in approach, the estimated robust optimal policies and worst case utilities are $C(\tilde{\theta}(P_{\theta_0,u}))$ and $V(\tilde{\theta}(P_{\theta_0,u}))$, respectively. In this way, we obtain a convenient description of the impact of a potential distortion u on estimated optimal policies and worst case utilities. Assuming a small model distortion u, we can then approximate the estimated robust optimal policy using a first order Taylor expansion around the true optimal robust policy $C(\theta_0)$ as follows:

$$C(\tilde{\theta}(P_{\theta_0,u})) - C(\theta_0) = \frac{\partial C(\theta_0)}{\partial \theta'} \frac{\partial \tilde{\theta}(P_{\theta_0})}{\partial P_{\theta_0,u}} + o(|P_{\theta_0,u} - P_{\theta_0}|) \ , \tag{14}$$

where $\frac{\partial \tilde{\theta}(P_{\theta_0})}{\partial P_{\theta_0,u}}$ is a (functional) Gâteaux derivative and $|\cdot|$ is a suitable norm on the vector space of finite signed measures. Therefore, to first order, the asymptotic sensitivity of estimated optimal policy $C(\tilde{\theta}(P_{\theta_0,u}))$ to distortion u is bounded if and only if the sensitivity of point estimate $\tilde{\theta}(P_{\theta_0,u})$ to distortion u is bounded. A similar result applies for the sensitivity of estimated worst case utility $V(\tilde{\theta}(P_{\theta_0,u}))$.

Overall, these arguments show that in order to limit the sensitivity of estimated robust optimal policies and worst case utilities to a potential model distortion u, it is sufficient to consider estimators having a bounded sensitivity with respect to such model misspecifications.

3.2. Bounded Utility Sensitivity and Bounded Influence Estimators

Huber's (1964) estimator is an example of so-called bounded influence (or, bounded influence function) estimators, which aim at bounding the sensitivity of the resulting point estimates to arbitrary local deviations from a given benchmark model. Introduced by Hampel (1974), the influence function is the first derivative of an estimator, viewed as a functional of the underlying DGP, and describes the linearized asymptotic bias of a statistic when adding a single contaminated observation.[8] Bounded influence estimators are estimators with a bounded influence function and imply a bounded sensitivity to arbitrary local deviations from a fixed reference model. Therefore, according to approximation (14) they also imply a bounded sensitivity of estimated robust optimal policies and worst case utilities $C(\tilde{\theta}(P_{\theta_0,u}))$ and $V(\tilde{\theta}(P_{\theta_0,u}))$ in general contexts of ambiguity aversion. In terms of the discussion in the previous section, the fact that P_{θ_0} is only an approximate description of the true DGP requires estimators with a bounded influence function, which is typically not the case for (pseudo) maximum likelihood estimators.

For practical purposes, a broad class of optimal robust estimators with bounded influence function is easily derived in a M, GMM or EMM type estimation setting, simply by replacing the unbounded M , GMM or EMM estimating function $\mathbf{s}(\cdot)$ by a weighted estimating function $\boldsymbol{\psi}_c(\mathbf{s}(\cdot))$; see, among others, Mancini, Ronchetti, and Trojani (2005) and Ortelli and Trojani (2005). Let $\mathbf{y}$ and $\boldsymbol{\theta}$ be random observations from the underlying DGP and the vector of parameters of interest, respectively. The robust bounded influence estimator $\widehat{\boldsymbol{\theta}}_{T,\mathrm{rob}}$ is the M, GMM or EMM type estimator based on orthogonality condition $\mathbb{E}\left[\boldsymbol{\psi}_c(\mathbf{s}(\mathbf{y},\boldsymbol{\theta}))\right] = \mathbf{0}$, where the Huber-weighted estimating function $\boldsymbol{\psi}_c(\mathbf{s}(\mathbf{y},\boldsymbol{\theta}))$ is given by:

$$\boldsymbol{\psi}_c(\mathbf{s}(\mathbf{y},\boldsymbol{\theta})) = \mathbf{A}(\boldsymbol{\theta})(\mathbf{s}(\mathbf{y},\boldsymbol{\theta}) - \boldsymbol{\tau}(\mathbf{y},\boldsymbol{\theta}))\mathbf{w}(\mathbf{y},\boldsymbol{\theta}), \tag{16}$$

---

[8] The influence function of a statistical functional $\tilde{\theta}$ at reference model P_{θ_0} and contamination direction x is formally defined as:

$$\mathrm{IF}_{\tilde{\theta}}(P_{\theta_0}, x) = \lim_{h \to 0} \frac{\tilde{\theta}\left((1-h)P_{\theta_0} + h\delta_x\right) - \tilde{\theta}\left(P_{\theta_0}\right)}{h} \ , \tag{15}$$

where δ_x is the dirac measure in x.

for some matrix $\mathbf{A}$, a bias correction vector $\boldsymbol{\tau}$ and Huber weights $\mathbf{w}$, defined for $c \geq \sqrt{\dim(\mathbf{s}(\cdot))}$ by:

$$\mathbf{w}(\mathbf{y}, \boldsymbol{\theta}) = \min\left(\mathbf{1}, c \left\|\mathbf{A}(\boldsymbol{\theta})(\mathbf{s}(\mathbf{y}, \boldsymbol{\theta}) - \boldsymbol{\tau}(\mathbf{y}, \boldsymbol{\theta}))\right\|^{-1}\right), \tag{17}$$

where $\|\cdot\|$ denotes the Euclidean norm. The weighted estimating function $\boldsymbol{\psi}_c(\mathbf{s}(\cdot))$ is bounded and therefore ensures a bounded sensitivity of estimator $\widehat{\boldsymbol{\theta}}_{T,\text{rob}}$ with respect to small distortions of the given reference model. The constant c controls the tradeoff between efficiency under the ideal model and robustness under a model misspecification: The lower c the more robust the estimator, at the cost of a potentially lower efficiency under ideal model conditions.[9]

Overall, robust M, GMM or EMM type estimators of the above type can be derived for a fairly general class of dynamics models, allowing to explore the link between ambiguity aversion and robust estimation in a large variety of dynamic specifications of models with ambiguity, featuring, e.g., non-Gaussian returns or state-dependent volatility and predictability features. In the following sections, we quantify in a more systematic way the main implications of our analysis for such general dynamic settings of ambiguity.

4. Taking Ambiguity to Reality

Many classical estimators, like, e.g., several PML type estimators, have an unbounded sensitivity to misspecifications of the benchmark model. Therefore, their in-sample point estimates can highly depend on the specific structure of unknown time-varying DGP features and be highly misleading for the implementation of out-of-sample optimal policies. In this case, a systematic approach to downweight in-sample observations that are likely too influential in presence of a model deviation is necessary to produce estimated optimal decision rules with a performance not largely dependent on the potential model distortion selected by nature. The general robust estimation approach of Section 3 provides a coherent framework to achieve this goal. Downweighting the most influential observations relative to, e.g., a maximum likelihood setting, can be interpreted as not trusting some of the data completely, by accounting for the fact that the maximum likelihood assumptions are likely not entirely satisfied. In this sense, a robust agent estimating a robust optimal policy cannot realistically trust all observations the same way, because some of them are potentially more damaging than others, in presence of a distortion of the ideal model, for the properties of estimated robust optimal policies.

---

[9]The matrix $A(\boldsymbol{\theta})$ in Equation (16) ensures that the scaling condition

$$\mathbb{E}_{\boldsymbol{\theta}_0}\left[\boldsymbol{\psi}_c(\mathbf{s}(\mathbf{y}, \boldsymbol{\theta}_0))\boldsymbol{\psi}_c(\mathbf{s}(\mathbf{y}, \boldsymbol{\theta}_0))^\top\right] = \mathbf{I} \tag{18}$$

is satisfied, implying that the norm of the self-standardized influence function of this estimator is bounded by c. Since the bias correction $\boldsymbol{\tau}$ in Equation (16) is rarely available in closed form, one would have to solve multi-dimensional integrals using Monte Carlo simulations, which can be a numerically daunting task. A convenient alternative to circumvent the computation of $\boldsymbol{\tau}$ all together is to correct the bias induced by the reweighting of unbounded estimating function $\mathbf{s}(\cdot)$ using the robust Efficient Method of Moments proposed by Ortelli and Trojani (2005). In this case, $\tau = 0$ and $A(\boldsymbol{\theta})$ is directly obtained from solving the equation:

$$\mathbf{A}(\boldsymbol{\theta})^T\mathbf{A}(\boldsymbol{\theta}) = \mathbb{E}\left[\mathbf{s}(\mathbf{y}, \boldsymbol{\theta})\mathbf{w}(\mathbf{y}, \boldsymbol{\theta})\left(\mathbf{s}(\mathbf{y}, \boldsymbol{\theta})\mathbf{w}(\mathbf{y}, \boldsymbol{\theta})\right)^T\right]^{-1}. \tag{19}$$

The updating algorithm for determining $\mathbf{A}$ is given in Appendix A.

These arguments are valid more generally, e.g., for an econometrician aiming at estimating the general equilibrium parameters of an economy with ambiguity aversion. By construction, all equilibrium quantities in such economies, including the market prices of risk and ambiguity, are determined by the parameters of a fixed reference model. Thus, they do not depend on the specific form of a potentially time-varying distortion in the actual DGP dynamics. In contrast, the econometrician observes asset price dynamics that are potentially distorted by such unknown time-varying features. In order to estimate general equilibrium parameters that do not heavily depend on the potential (time-varying) deviation from the reference belief in the general equilibrium economy, she might therefore need to downweight the effects of influential observations, relative to a context where time-varying model misspecifications are assumed a priori as not existent. Also in these settings, the robust tools introduced in Section 3 offer a coherent framework to achieve this objective.

In the following sections, we show how ambiguity can be coherently taken to reality, by applying the robust estimation approach of Section 3. To demonstrate the broad applicability of our approach to different settings of ambiguity, we consider in more detail three distinct applications. First, we study a robust portfolio choice problem with event risk. Second, we estimate the risk and ambiguity premium in a jump-diffusion general equilibrium economy. Finally, we estimate bid and ask option prices in a market economy with ambiguity.

4.1. Robust Decision Rules and Estimation with Event Risk

Unknown time-varying features affecting the DGP of financial variables can be related to different aspects of the conditional distribution, including, e.g., skewness, kurtosis and event risk features. As noted by Liu, Pan, and Wang (2005), event risk characteristics are very difficult to estimate precisely, because they are related to events that are rarely observed in financial markets. Similarly, higher order moments like conditional skewness and kurtosis are more difficult to estimate than conditional means and variances, because they largely depend on the properties of rare tail events. Therefore, unknown time-varying DGP features can have dramatic effects on optimal policies that depend on ideal assumptions about the probabilistic structure of rare event risk, making ambiguity aversion an even more natural modeling assumption in these settings. In order to study the link between ambiguity aversion and robust estimation in this framework, we consider an ambiguity averse agent who selects robust optimal portfolios subject to continuous-time dynamics with jumps in returns. In this context, we quantify the tradeoff between ambiguity aversion and robust estimation for the out-of-sample utility of estimated optimal policies.

4.1.1. Robust Portfolio Choice with Event Risk

We consider a simple Merton (1976) jump-diffusion (JD) as a reference model for asset returns:

$$dS(t) = \mu_S S(t)dt + \sigma S(t)dB(t) + \left(e^{\xi^Y} - 1\right) S(t)dN_t, \tag{20}$$

with $\mu_S = \mu + \frac{1}{2}\sigma^2$, an iid jump size $\xi^Y \sim N(\mu_Y, \sigma_Y^2)$ and a Poisson process N with constant intensity λ_Y. While this assumed reference model is more realistic for returns, compared, e.g., to a simple GBM, the parameters of the jump component are inherently difficult to estimate and might contain small time-varying features, which could be even harder to identify using realistic sample sizes. Therefore, the agent exhibits ambiguity aversion with respect to rare events and wants to protect herself against these unreliable aspects of her reference model.

To model ambiguity aversion, we follow Liu, Pan, and Wang (2005) and assume an agent who derives robust portfolio rules for investing a share π_t into the jump-diffusion stock and the remainder at the risk free rate. The ambiguity agent then faces the following robust portfolio choice problem:

$$V(W_t, t) = \sup_{\{C_t, \pi_t\}} \inf_{\{a,b\}} E^{\{a,b\}} \left[\int_0^T e^{-\beta t} \left(\frac{C_t^{1-\gamma}}{1-\gamma} + \frac{1}{\Psi} H(a,b) \right) dt \right], \tag{21}$$

subject to:

$$dW_t = \left[W_t \left(r + \pi_t \left(\mu_S - r \right) \right) - C_t \right] dt + \pi_t \sigma W_t dB_t + \pi_{t-} W_{t-} (e^{\xi^Y(b)} - 1) dN_t(a), \tag{22}$$

where a and b are parameters that describe a potential deviation from the reference model, and Ψ is a state dependent control for the level of robustness sought. Time preference rate and risk aversion are given by β and γ, respectively. Function $H(\cdot)$ measures the discrepancy between reference and alternative models, parameterized by (a, b). It is given by:

$$H(a,b) = \lambda \left[1 + \left(a + \frac{1}{2} b^2 \sigma_Y^2 - 1 \right) e^a + d(1 + (e^{a + b^2 \sigma_Y^2} - 2) e^a) \right] \tag{23}$$

for some $d > 0$.[10] Alternative models for the jump component are specified by appealing to the Radon-Nikodym derivative for jump size and intensity:

$$dZ_t = \left(e^{a + b\xi_t^Y - b\mu_Y - \frac{1}{2} b^2 \sigma_Y^2} - 1 \right) Z_{t-} dN_{t-} - \left(e^a - 1 \right) \lambda_Y Z_t dt. \tag{24}$$

This Radon-Nikodym process changes the jump intensity from λ_Y under the reference model to $\lambda_Y^Z = \lambda_Y e^a$ for the considered alternative. Similarly, the mean jump size $E\left[e^{\xi^Y}\right] = e^{\mu_Y + \frac{1}{2}\sigma_Y^2}$ under the reference model is transformed to a mean jump size $E^{\{b\}}\left[e^{\xi^Y}\right] = e^{\mu_Y + \frac{1}{2}\sigma_Y^2} \cdot e^{b\sigma_Y^2}$ under the distorted alternative. The supplemental appendix derives the solution for the CRRA robust control problem when $\Psi(W,t) = \frac{\vartheta}{(1-\gamma)V(W,t)}$, where as in previous sections parameter ϑ controls the degree of ambiguity aversion.

4.1.2. Modeling Unknown Time-Varying Event Risk Features

To quantify the potential implications of robust estimation in the context of the previous subsection, we assume that returns over a period of $H = 6$ years follow a small time-varying

[10] $H(\cdot)$ is the continuous time limit of an extended entropy measure; see Liu, Pan, and Wang (2005) for details.

perturbation of the JD reference model in Equation (20):

$$dS(t) = \mu_S S(t)dt + \sigma S(t)dB(t) + \left(e^{\chi_t^Y} - 1\right) S(t)dN_t \quad ; \quad t \in [0, H], \tag{25}$$

where χ_t^Y is an iid jump size drawn from a time-varying distribution $N(\mu_{Y,t}, \sigma_Y^2)$ and $\mu_{Y,t}$ is a piecewise constant function of time given by:

$$\mu_{Y,t} = \begin{cases} -0.04 & t \in [0, 0.21H) \\ -0.01 & t \in [0.21H, 0.54H) \\ 0.01 & t \in [0.54H, 0.58H) \\ 0.12 & t \in [0.58H, 0.75H) \\ -0.02 & t \in [0.75H, 0.83H) \\ -0.1 & t \in [0.83H, H]. \end{cases} \tag{26}$$

The unconditional average jump size implied by the dynamics in Equation (25) is $\mu_Y = -1\%$, as in agent's reference model, but it conditionally oscillates around this value. For instance, in the first 1.25 years the average jump size is -3%, instead of -1%. In general, this type of distortion is small: Jumps occur infrequently, and certain jumps may both be compatible with conditional and unconditional jump size distributions. Due to the symmetric structure of the distortion, these time-varying jump size features alter the tail behavior of returns, but not their unconditional mean. The top left panel of Figure 2 illustrates the density of returns in the distorted model (Equation (25)), together with the density of a reference JD model with parameters $\mu = 0.08, \sigma = 0.15$, $\lambda_Y = 3$, $\mu_Y = -0.01$, $\sigma_Y = 0.04$, and suggests a very small discrepancy between these two distributions: As in the GBM case, even if the tail behavior of the two distributions is different, without good prior knowledge, it is extremely difficult to detect the specific structure of the time-varying jump component in Equation (25).

[Figure 2 about here.]

A sample realization for the distorted process (top right panel) and the decomposition into clean data (bottom right panel) and time-varying jump component (bottom left panel) further highlight the difficulty to distinguish the two processes based on realistic sample sizes of returns.

4.1.3. Ambiguity Aversion, Robust Portfolio Rules and Estimation

Unless parameters are common knowledge, the ambiguity averse agent needs to obtain estimates for her reference model, before she can implement her (robust) portfolio policies. The analytical expression for the likelihood function of the reference model in Equation (20) is provided in Appendix B. The computation of the first order conditions implied by this likelihood function yields an unbounded estimating function $\mathbf{s}(\mathbf{y}, \boldsymbol{\theta})$, which implies the non-robustness of maximum likelihood estimators in this context. In order to derive a robust estimator for the given setting, we apply the general robust estimation approach outlined in Section 3, and derive robust estimators for JD Model (20) using the weighted estimating function implied by Equation (16).

To quantify the tradeoff between ambiguity aversion and robust estimation in the jump-diffusion setting, we consider for estimation a sample of five years of daily market returns generated by the distorted model (Equation (25)). We then compute the one year out-of-sample utility of the estimated robust optimal policies. Panel (b) of Table 3 presents the corresponding out-of-sample realized utility, in wealth equivalents, for (i) the case of known model parameters, (ii) parameters estimated by maximum likelihood and (iii) parameters estimated by the robust estimator. For comparison, Panel (a) shows the implied out-of-sample utility in the case where both in- and out-of-sample returns are generated according to the agent's reference model.

[Table 3 about here.]

When the return generating process is a clean jump-diffusion and model parameters are common knowledge (first row of Panel (a)), the expected utility agent ($\vartheta = 0$) attains, as expected, the highest degree of out-of-sample utility. However, the loss in out-of-sample utility of ambiguity averse agents ($\vartheta > 0$) in this case is virtually non-existent and less than 0.2%: The price the ambiguity averse agent needs to pay in this ideal setup in order to be sheltered from a deviation from the reference model is fairly small. When relaxing the assumption of known parameters, agents have to pay an additional price, in terms of realized wealth, for parameter estimation. While this is generally not avoidable, two interesting aspects are worth mentioning. First, ambiguity aversion results in the agent having to suffer fewer wealth losses (less than 1%), compared to the expected utility maximizer (roughly 2%). Second, robust estimators (third row of Panel (a)) do not significantly reduce utility, when compared to maximum likelihood estimators (second row of Panel (a)), despite the latter being optimal in this setup: Overall, the costs of using robust estimators under ideal model assumptions is negligible in this jump-diffusion setting.

As expected, when returns follow the perturbed JD (25), the agent's utility is generally lower than in the clean jump-diffusion case. When parameters are common knowledge (first row of Panel (b)), a concern for misspecification of the jump component now improves the agent's utility. While this is expected and the reason for ambiguity averse choices, the additional utility impact deriving from parameter estimation is striking. Panel(b) of Table 3 shows that the difficulty of identifying rare event components significantly reduces the agent's utility deriving from classical estimation techniques. For instance, an expected utility maximizer loses more than one third of the wealth, compared to the known parameter case, when using maximum likelihood estimators; Robust portfolio rules help to recover only part of these utility losses. In contrast, robust estimators allow the agent to obtain out-of-sample utility losses that are uniformly smaller: The estimation induced utility loss of robust compared to maximum likelihood estimators is smaller by a factor of ten; Moreover, the utility losses relative to the known parameter case are less than about 3%, independent of the degree of ambiguity aversions. Finally, it is important to note that, similar to the previous results, also in this jump-diffusion setting the excessive sensitivity of optimal policies estimated by PML is an intrinsic feature of policies estimated by non-robust PML methods. Indeed, reversing the direction of time-varying mean jump size contamination relative to the unconditional jump size of -1% in Equation (25), also results in the robust estimator outperforming PML in terms of realized out-of-sample utility. Table I.1 of the supplemental appendix provides detailed results.

Overall, these results indicate that the robust estimation approach of Section 3 successfully shelters the agent from the potential damaging effects of unknown time-varying components in the data generating process featuring event risk. In contrast, non-robust estimators can imply very dramatic losses, already under very small time-varying deviations from the ideal JD assumptions.

4.2. Event Risk and the Equilibrium Market Prices of Risk and Ambiguity

In equilibrium, event risk ambiguity has to be priced and reflected by the time series and the cross-sectional behavior of asset prices. Liu, Pan, and Wang (2005) show that ambiguity aversion about event risk can help to reconcile the equity premium and some of the smirk patterns of index options. Drechsler (2009) proposes a general equilibrium economy with long run risk, rare disasters and time-varying ambiguity aversion that captures the variance risk premium and option skew, while simultaneously matching the moments of cash-flows and stock returns.

Given the specification of a general equilibrium economy with ambiguity, the relevant equilibrium quantities are, by construction, independent of the form of candidate (potentially time-varying) distortions of the reference model dynamics for asset prices. However, the econometrician observes asset prices in real time, from data that have been potentially distorted by such unknown time-varying features. Therefore, to estimate equilibrium parameters, she might have to be cautious and not use methods and point estimates that heavily depend on the structure of a (time-varying) deviation from agents' reference model. This simple intuition motivates the usefulness of the robust estimation approach introduced in Section 3 also for the estimation of the parameters of general equilibrium economies with ambiguity. To understand this general equilibrium induced link between ambiguity aversion and robust estimation in a simple setting, we consider a general equilibrium version of the partial equilibrium rare event economy introduced above and take the point of view of an econometrician who wants to estimate ambiguity and risk premiums based on time series of option implied volatility smiles and the corresponding underlying returns.

4.2.1. The Cross Section of Derivative Prices: Identifying the Market Price of Event Risk and Ambiguity

Studying the general equilibrium problem of a representative agent with preferences given by Equation (21) allows for the computation of equilibrium premiums for risk and ambiguity, as well as equilibrium option prices. Given the market clearing condition $\pi_t^* = 1$, the supplemental appendix derives the implicit solutions for the equilibrium worst case choices a^* and b^*, which together with the representative agent's first order condition for optimal portfolio choice imply the following equity premium in the jump-diffusion economy:

$$\mu_S - r \quad = \quad \gamma\sigma^2 - \lambda_Y e^{a^*}\left(e^{\mu_Y(1-\gamma)+\frac{1}{2}(1+\gamma^2)\sigma_Y^2(1+2b^*)} - e^{-\gamma\mu_Y+\frac{1}{2}\sigma_Y^2\gamma^2(1+2b^*)}\right). \tag{27}$$

How does this equity premium compare to that of standard economies? Clearly, the Black-Scholes (BS) environment implies $\lambda_Y = 0$, while the jump-diffusion (JD) setting without ambiguity emerges for $a^* = b^* = 0$:

$$\textbf{BS:} \qquad \gamma\sigma^2, \tag{28}$$

$$\textbf{JD:} \qquad \gamma\sigma^2 - \lambda_Y\left(e^{\mu_Y(1-\gamma)+\frac{1}{2}(1+\gamma^2)\sigma_Y^2} - e^{-\gamma\mu_Y+\frac{1}{2}\sigma_Y^2\gamma^2}\right). \tag{29}$$

Given a choice of parameters for the reference belief, Table 4 shows the relation between the equity premium and the preference parameters governing risk and ambiguity aversion.

[Table 4 about here.]

For instance, for $\gamma = 3$ we see that moving from a Black-Scholes to a jump-diffusion economy nearly doubles the equity premium from 6.75% to 11.08%, under the given parameter choice. When the representative agent is additionally ambiguity averse (i.e. $\vartheta > 0$), the equity premium increases by approximately another 4% for each increase in parameter ϑ.

Given the market incompleteness arising from the presence of jumps in the asset return process, the econometrician needs additional information, e.g., from a time series of option prices written on the underlying JD equity index, to decompose the latent structure of the equity premium into a risk and an ambiguity premium component. The supplemental appendix shows that equilibrium option prices have to satisfy the following fundamental differential equation:

$$\begin{aligned} rC_t = \frac{\partial C}{\partial t} &+ \left(r - \lambda_Y e^{a^*} E^{\{b^*\}}\left[e^{(1-\gamma)\xi^Y} - e^{-\gamma\xi^Y}\right]\right) S_t \frac{\partial C}{\partial S} + \frac{1}{2}\frac{\partial^2 C}{\partial S^2}\sigma^2 S_t^2 \\ &+ \lambda_Y e^{a^*} E^{\{b^*\}}\left[e^{-\gamma\xi^Y}\left(C_t - C_{t-}\right)\right], \end{aligned} \tag{30}$$

subject to the standard boundary conditions. The solution of this equation provides the equilibrium pricing function $C_t = C(S_t, \boldsymbol{\theta}, \gamma, \vartheta)$ in the economy with ambiguity aversion, where parameter $\boldsymbol{\theta} = (\mu, \sigma, a, \sigma_Y)'$ collects the parameters describing the stock return dynamics under the representative agent reference belief. Note that this pricing function depends only on reference belief and representative agent preference parameters. Therefore, the cross section of option prices is independent on the emergence of potentially time-varying unknown features in the dynamics of underlying returns. Figure 3 visualizes the dependence of the equilibrium option implied volatility smile on the representative agent's risk and ambiguity aversion parameters in the economy. As highlighted in Liu, Pan, and Wang (2005), risk and ambiguity aversion impact on both the level and the slope of the smile, allowing the econometrician to use option-implied information in order to identify risk and ambiguity premiums in this setting.

[Figure 3 about here.]

4.2.2. Market Prices of Risk and Ambiguity: Robust Estimation

We assume an econometrician that observes the equilibrium interest rate r, a time series of index prices $S_1, \ldots, S_T$ and a time series of equilibrium option prices $C(\boldsymbol{\theta}, S_1), \ldots, C(\boldsymbol{\theta}, S_T)$. Using this information, she estimates the general equilibrium parameters and, in particular, the equilibrium risk and ambiguity premium in the economy. This task is complicated by the fact that the reference model dynamics in Equation (20) are potentially distorted by unknown time-varying features. Note that since the representative agent is ambiguity averse, the specific features of such potential model distortions do not appear in any of the equilibrium parameters of interest. Therefore, it is intuitively clear that they should also not strongly influence the point estimates of these parameters produced by a robust estimation approach consistent with the structure of the general equilibrium economy with ambiguity aversion.

Figure 4 plots the finite sample distributions of estimated equity premiums, as well as those of the corresponding risk and ambiguity premiums, obtained when using (i) maximum likelihood estimators or (ii) the robust estimation approach from Section 3. In order to decompose the equity premium in its components, we assume that option prices are observed without error by the econometrician, so that risk and ambiguity premiums are simply estimated by inverting the equilibrium option pricing function.[11]

[Figure 4 about here.]

The right panels of Figure 4 present results for a return DGP that exactly follows the reference model (Equation (20)). The panels on the left present results for the case where the DGP is distorted by unknown time-varying features according to Equation (25). The estimation with clean data shows that risk and ambiguity premiums are estimated without any significant bias by all methods. Moreover, the finite sample distribution of premiums estimated by classical and robust methods is virtually indistinguishable, confirming that the efficiency costs of using robust instead of maximum likelihood methods in the ideal JD setting are negligible. The situation is dramatically different in presence of a time-varying distortion: When the data contain small time-varying components, Panel (a) shows an equity premium estimated by maximum likelihood that is systematically underestimated. While robust methods imply an unbiased estimation of the true equity premium of 8.26%, the classical estimator yields a premium that is over nine percentage points smaller on average. Therefore, non-robust estimation approaches can lead to significant misperceptions of the size of the equity premium in presence of even small time-varying distortions of the benchmark model. Panels (c) and (e) further decompose the estimated equity premium into the risk and the ambiguity premium, respectively. We find that under time-varying distortion (Equation (25)), the ambiguity premium can be consistently estimated at around 0.4% by both estimators, while it is the (diffusive and jump) equilibrium

[11] For maximum likelihood estimation, we use the analytical expression for the likelihood function of the reference JD model (Equation 20) provided in Appendix B. We derive robust estimators for the JD model using the weighted estimating function implied by Equation (16). Given the estimated JD parameters, we then invert the equilibrium option pricing function to extract the implied risk and ambiguity premium components.

risk compensation that is underestimated by non-robust estimators.[12]

Overall, these findings confirm that our robust estimation approach can also consistently estimate the general equilibrium structure of economies with ambiguity aversion when small unknown time-varying features are present in the DGP of observed asset returns.

4.3. Option Pricing under Ambiguity

While previous sections have focused on utility specifications for decision making, the outlined econometric tools for robustification are explicitly not restricted to estimating utility models of ambiguity aversion. This section illustrates this point by considering option pricing in incomplete markets. More precisely, assume that an agent wants to quote option prices in a setting where she is uncertain about the true DGP. As in the previous section, her reference model is a jump-diffusion. Even without ambiguity about the jump component, the market is incomplete and there exist infinitely many arbitrage free prices for the option unless prices of traded options are available. The presence of ambiguity introduces an additional degree of incompleteness, because the agent considers alternative models that imply different market prices of risk. To narrow down the set of potential prices, a number of papers have proposed to not only exclude prices which violate the no arbitrage condition, but also prices which are too favorable in terms of sharp ratio (see, e.g., Cochrane and Saá-Requejo (2000) or Björk and Slinko (2006)). Along these lines, bid and ask prices in an ambiguity setting can be found by deriving minimally and maximally acceptable option prices given some level of potential model misspecification.

For simplicity and to isolate the effect of ambiguity, assume that the market price of jump risk is known and equal to zero, then pricing bounds for a given maximum entropy η between reference model and alternatives are given by:

$$\inf_{a,b} C^{(a,b)}(S,K) \qquad \text{subject to} \quad H(a,b) \leq \eta \tag{31}$$

and

$$\sup_{a,b} C^{(a,b)}(S,K) \qquad \text{subject to} \quad H(a,b) \leq \eta, \tag{32}$$

where $C(\cdot)$ is the price of the call option with strike K and initial stock price S. This option price satisfies the following partial differential equation:

$$\begin{aligned} rC_t = \frac{\partial C}{\partial t} + \left(r - \lambda_Y e^{a^*} E^{\{b^*\}}\left[e^{\xi^Y} - 1\right]\right) S_t \frac{\partial C}{\partial S} + \frac{1}{2}\frac{\partial^2 C}{\partial S^2}\sigma^2 S_t^2 \\ + \lambda_Y e^{a^*} E^{\{b^*\}}\left[C_t - C_{t-}\right], \end{aligned} \tag{33}$$

subject to the usual boundary conditions. The benchmark price is obtained by setting $a = b = 0$ giving the reference model Merton price for which the closed form solution is well known; $H(a,b)$ is the relative entropy between reference model and alternatives and is defined by Equation (23).

[12] Similar to the exercise in Section 2.6, only robust estimates result in a consistent estimation of the equity premium when the contamination is reversed; results are available on request.

Figure 5 depicts bid and ask prices for different levels of moneyness (Subplot (a)) and ambiguity (Subplot (b)) when the parameter of the reference model are known. The presence of ambiguity about the true DGP leads to significant bid-ask spreads. Interestingly, model uncertainty causes the ask price to deviate more in absolute terms from the Merton price than the bid price. This can be related to especially the writing of a call requiring an additional compensation for uncertainty about the hedge position.

[Figure 5 about here.]

Once parameters need to be estimated, the sensitivity of classical estimates to time-varying features carries over to the pricing bounds. Subplots (a), (c), and (e) of Figure 6 show that robust parameters result in option prices being very close to the true price. For classical estimates, the true price is only included when considering 95% of the estimated price distribution (77% for the bid price). Also the dispersion of the estimated distributions is on average 37% smaller for robust estimates based on the inter-quartile range. Subplots (b), (d), and (f) show bid, reference model, and ask prices in absence of time-varying unknown data features, for comparison; here robust and classical parameter estimates perform equally well and estimated bid and ask prices are essentially identical to the true values.

[Figure 6 about here.]

The picture of the superior performance of robust estimates is confirmed when looking at Table 6 which shows (relative) bid-ask spreads. These are significantly larger for non-robust estimates and further deviate from the bid-ask spreads based on true reference model parameters, which has important economic implications. Investors who trade options are required to maintain cash reserves to support a trade. In line with Carr, Madan, and Alvarez (2010), these capital reserves need to be larger, the larger are bid-ask spreads. Intuitively, an investor who buys the option at the ask, but has to sell it again at the bid price incurs cost equal to the bid-ask spread. Therefore, the investor needs to hold enough capital to be able to cover the cost of such an unfavorable unwinding of the trade.

Interpreting the results of Table 6 in these economic terms, robust parameter estimates significantly lower capital requirements. For out-of-the money options, the capital charge based on classical estimates can be four times larger than the capital required when robustly estimating the parameters of the reference model. Even for in-the-money options the investors who uses classical estimates is required to hold at least 50% more capital. Consequently, costly capital reserves are minimized when using robust parameter estimates.

[Table 6 about here.]

5. Empirical Application: Ambiguous Predictability and Robust Estimation

Motivated by the results and intuitions of the previous sections, where we analyzed the added economic value of our robust estimation approach in different economies with ambiguity aversion, theoretically and by Monte Carlo simulation, this section studies a real data empirical application to stock return predictability. We apply the robust approach developed in Section 3 to define robust estimators for ambiguous predictability structures in US stock return data, and we quantify the resulting economic implications for (robust) asset allocation.

5.1. Unknown Time-Varying Components and Ambiguous Predictability Features

A large literature studies the predictability of stock returns by macroeconomic or financial indicators, such as the dividend-price ratio. Theoretically, in an arbitrage free-market, time-variation of, e.g., dividend-price ratios is already a hint of potential predictability structures, either for future dividends or stock returns. The empirical literature tends to find more evidence of stock return predictability, especially at quarterly or yearly horizons. However, the strength of this evidence is mixed, partly because it significantly depends on the choice of the sample under investigation, and it is difficult to exploit estimated predictability features to produce additional economic value, e.g., in the context of real-time asset allocation; see Lettau and Ludvigson (2010) and Welch and Goyal (2008), among others, for an excellent review of this literature.

The difficulty in identifying predictability structures can be enhanced by the presence of time-varying features in the underlying data generating process. Lettau and Van Nieuwerburgh (2008), for instance, suggest that a permanent downward shift in the mean of the log dividend-price ratio during the mid-1990s can help to reconcile some of the mixed results on stock return predictability. However, as noted by Lettau and Ludvigson (2010), such a structural break hypothesis is hardly consistent with the recent reversal to higher average dividend-price ratios, suggesting that the mid-1990s evidence was more likely associated with a temporary unusual period, rather than with a structural shift in the overall structure of aggregate financial ratios. From a broader perspective, the observed pattern in log dividend-price ratios is thus potentially consistent with the presence of unknown time-varying components in the data generating process, following the intuition provided by our previous findings.

Interestingly, the empirical distribution of yearly log dividend-price ratios exhibits excess kurtosis that can be consistent with data generated by slightly distorted models. Consider, for instance, the standard predictability setting, in which returns r_t are forecasted by a persistent dividend-price ratio x_t, modeled as an autoregressive process:

$$\begin{aligned} r_{t+1} &= \delta + \beta x_t + \varepsilon_{1,t+1}, \\ x_{t+1} &= \omega + \varphi x_t + \varepsilon_{2,t+1}, \end{aligned} \tag{34}$$

with normal innovations $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ having correlation ϱ. Given the persistence of the predictor variable, a perturbation to ω at a single point in time might generate a pattern similar to the one observed empirically, because one unusual realization can impact the process over a long period. This feature is illustrated in Figure 7, which compares observed data (Panels (a), (c), and (e)) to simulated sample paths for excess stock returns and log dividend-price ratios (Panels (b), (d), and (f)). The parameters used for the simulation are estimated from annual data based on

Center for Research in Security Prices (CRSP) value-weighted market returns during 1926–1994, which excludes the potential structural break period described in Lettau and Van Nieuwerburgh (2008).

[Figure 7 about here.]

The dashed lines in Panel (b) and (d) show the sample paths for the clean process in Equations (34) and (34): This log dividend-price process is rather stable and does not include the significant drop around the year 2000. The simulated trajectory underlying the solid lines is identical, except that the clean process is perturbed by an unknown time-varying component. More precisely, replacing $\omega = -0.40$ by $\tilde{\omega} = -1.25$ at a single point in time during 1997 is sufficient to generate a pattern that quite closely resembles the real data. Panels (e) and (f) confirm these findings, showing that skewness and kurtosis of the empirical distribution of the distorted log dividend-price series roughly match those observed for real data. Consequently, a data generating process including unknown time-varying components can be a valid alternative hypothesis for the observed pattern in yearly return and dividend-price ratio data, motivating an ambiguous modeling of predictability features.

For higher, e.g., monthly, data frequencies, the additional data information can improve the efficiency of parameter estimates, but a possibly more complex structure of unknown time-varying features can make the detection of ambiguous predictability structures even more difficult. To illustrate this aspect more concretely, Figure 8 presents dividend-price ratios and excess returns of the CRSP value-weighted market portfolio, together with an arbitrary simulated sample path of (i) a monthly Gaussian VAR model (Equation (34)) and (ii) a Gaussian VAR model contaminated by a time-varying jump component: The inclusion of the jump component in the stock price process tries to account for potential unknown time-varying features not captured by the simple Gaussian VAR assumption.

[Figure 8 about here.]

As can be seen from a simple comparison of the simulated red solid line in Panel (d) and the actual return series in Panel (c), the inclusion of the jump component helps to better reconcile the stochastic feature of the data generated from the Gaussian VAR model with those generated by the actual DGP. Thus, a Gaussian VAR model contaminated by unknown time-varying components is a potentially valid alternative hypothesis to a Gaussian VAR for monthly returns and dividend-price ratios, motivating an ambiguous modeling of predictability features also at the monthly frequency.

It is useful to illustrate the potential impact of unknown time-varying features on estimated predictability structures at monthly horizons. Figure 9 shows the empirical distributions of estimated predictability and correlation coefficients in a Monte Carlo simulation of the distorted Gaussian VAR model.

[Figure 9 about here.]

We estimate the VAR parameters using (i) the maximum likelihood estimator implied by the Gaussian VAR assumption and (ii) the associated robust maximum likelihood estimator implied by the bounded influence methodology introduced in Section 3. Figure 9 indicates that predictability parameters estimated by the classical estimator under the distorted VAR model are less efficient than those produced by the robust estimation methodology. At the same time, the negative correlation between returns and dividend-price ratios estimated by classical methods is significantly upward biased, relative to the one implied by the robust estimator. The lower efficiency for the predictability parameter and the positive bias for the correlation coefficient can have economically relevant implication in the construction of optimal dynamic portfolios and intertemporal hedging demands under ambiguous predictability structures. The following sections study the relation between robust estimation and dynamic portfolio choice with ambiguous predictability in more detail.

5.2. Ambiguous Predictability and Robust Estimation

We estimate a standard predictive VAR model using monthly value-weighted stock index (with and without dividends) and T-Bill returns, obtained from CRSP, for the period 1929–2009. We construct continuously compounded excess market returns and dividend-price ratios following Fama and French (1988)[13]. Let the sampling frequency be $\Delta = 1/12$, i.e., one month, x_t be the log dividend-price ratio, and $r_{t+\Delta} = \log S_{t+\Delta} - \log S_t - \Delta r^f$ be the market index excess return over the (annualized) T-Bill risk free rate r^f. The corresponding predictive system reads:

$$\begin{bmatrix} r_{t+\Delta} \\ x_{t+\Delta} \end{bmatrix} = \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + \begin{bmatrix} 0 & d_3 \\ 0 & d_4 \end{bmatrix} \cdot \begin{bmatrix} r_t \\ x_t \end{bmatrix} + \begin{bmatrix} v_{1,t+\Delta} \\ v_{2,t+\Delta} \end{bmatrix}. \tag{35}$$

Parameter d_3 captures potential stock return predictability features, while parameter d_4 measures the persistence of dividend-price ratios. $v_t = (v_{1,t}, v_{2,t})'$ defines a bivariate martingale difference process with finite second moments. Given the potential ambiguity affecting specification (35), we estimate the predictive VAR using (i) the standard maximum likelihood estimator based on a Gaussian assumption for v_t and (ii) the corresponding robust maximum likelihood estimator implied by the robust methodology of Section 3. Figure 10 presents rolling estimates for predictive parameter d_3 implied by the two methods, with confidence interval bounds of plus and minus one standard error around each point estimate.[14]

[13]Dividend payments of the firms in the stock index are extracted as follows: In month t, the value of one dollar invested without reinvestment of dividends is $P(t) = \exp(r_{ND}(0) + r_{ND}(1) + \ldots + r_{ND}(t))$, where $r_{ND}(t)$ denotes the continuously compounded return of the stock index without reinvestment of dividends. Then, dividend payments in month t can be computed as $D(t) = P(t-1)\exp(r_D(t)) - P(t)$, where $r_D(t)$ is the continuously compounded return of the stock index with dividends reinvested. Given the strong seasonality in dividends paid by the companies in the stock index, the dividend-price ratio in month t is defined as the sum of the dividends paid in months $t-11$ to t, divided by the value of the stock index at time t.

[14]In every month t, we estimate the model based on sample data from January 1929 up to the current month. Classical standard errors are not adjusted for potential small sample biases, which might imply an even weaker evidence in favor of predictability using classical estimators and tests; see Campbell and Yogo (2006) or Amihud, Hurvich, and Whang (2009), among others.

[Figure 10 about here.]

The robust point estimate for d_3 is systematically larger, suggesting that bounding the effects of potentially influential data points helps to uncover a positive, even if time-varying, predictive relation between the dividend-price ratio and stock returns: While there is hardly any statistical evidence for predictability using the classical estimator, the robust estimator indicates the presence of predictability for certain sample periods. In line with Chen (2009), stock returns become predictable mainly in the postwar period with weaker evidence for predictability in recent years. Estimated parameters using the full sample of data are presented in Panel (a) of Table 7.

[Table 7 about here.]

In addition to the larger predictability parameter, robust estimates indicate a stronger negative correlation between the VAR innovations, a smaller variance of the error terms and a larger intercept in the return equation. These systematic differences between classical and robust point estimates suggest that indeed the findings provided by classical estimators might be highly dependent on the features of a limited fraction of observations. The robust methodology provides a simple way to verify this hypothesis, by inspecting the fraction of data for which a Huber weight less than one has identified an influential observation. Figure 11 plots excess returns and dividend yields over time, together with the corresponding Huber weight of each observation.

[Figure 11 about here.]

Overall, a fraction of 7.7% of the observations has been identified as particularly influential for classical parameter estimates. An important fraction of influential observations clusters during the Great Depression. In the postwar period, the stock market crash in 1987 has the largest influence, followed by observations having patterns similar to those of a sequence of isolated outliers. In particular, we see that influential observations often emerge in concomitance with large negative market returns. These patterns of influential observations nicely match the effects of ambiguity studied in the previous subsection, within a Gaussian VAR model contaminated by a time-varying jump component; see again Figure 8.

All in all, there appears to be some support in favor of return predictability with the Gaussian VAR model, at least when using the robust estimation approach. At the same time, the evidence of the presence of potentially influential observations suggests to treat the strict Gaussian VAR dynamics as ambiguous. The following sections study the economic implications of our robust estimation approach for optimal dynamic portfolio choice with ambiguous predictability features.

5.3. Ambiguous Predictability and Robust Portfolio Choice

We consider a robust version of the dynamic portfolio problem studied in Xia (2001). The reference model for the stock return dynamics is:

$$dS_t = \mu(X,t)S_t dt + \sigma S_t dB_t^{(1)}, \tag{36}$$

where $\sigma > 0$, $B^{(1)}$ is a standard Brownian motion, and the instantaneous expected return μ is related to a predictor variable X (the log dividend-price ratio):

$$\mu(X,t) = \alpha + \nu X_t, \tag{37}$$

with $\alpha, \nu \in \mathbb{R}$. Under the reference model, the predictor variable follows a persistent Gaussian process with dynamics:

$$dX_t = \kappa_X(\bar{x} - X_t)dt + bdB_t^{(2)}, \tag{38}$$

where $\kappa_X, \bar{x}, b > 0$ and $B^{(2)}$ is a second standard Brownian motion such that $dB_t^{(1)}dB_t^{(2)} = \rho dt$. This reference model is consistent with the discrete time model of Barberis (2000), who assumes the dividend yield, as predictor, to follow an AR(1) process. Investors take the reference model dynamics as ambiguous and solve a corresponding robust asset allocation problem:

$$V(t,W,X) = \sup_{\{\pi_t\}} \inf_{\{a\}} E^a \left[\frac{W_T^{1-\gamma}}{1-\gamma} + \int_0^T \frac{1}{2\Psi}(\pi_t \sigma W_t a)^2 dt \right], \tag{39}$$

subject to the dynamic budget constraint:

$$dW_t = \left[W_t \left(r + \pi_t \left([\bar{\mu} + \nu(X_t - \bar{x})] - r \right) \right) + \pi_t^2 \sigma^2 W_t^2 a \right] dt + \pi_t \sigma W_t dB_t^{(1)},$$

where $\bar{\mu} = \alpha + \nu\bar{x}$ and $\Psi = \vartheta/((1-\gamma)V(t,W,X))$, for $\gamma, \vartheta \geq 0$. This specification of ambiguity aversion in robust portfolio choice is borrowed from Maenhout (2006). The optimal portfolio rule π_t^* for an ambiguity averse investor with power utility follows similarly to the results of Section 4 and is given by:

$$\pi_t^* = \frac{1}{\gamma + \vartheta} \left(\frac{\mu(X,t) - r}{\sigma^2} + \left(B(\tau) + C(\tau) \left(\mu(X,t) - r \right) \right) \frac{\nu \rho b}{\sigma} \right), \tag{40}$$

where functions $B(\cdot)$ and $C(\cdot)$ are solutions to a system of ordinary differential equations, given in detail in the supplemental appendix, and $\tau = T - t$ is the investment horizon.

Parameters of the continuous time reference model in Equations (36) and (38) are easily estimated using its exact discretization. The supplemental appendix shows that Model (35) is the exact discrete time equivalent of reference model in Equations (36) and (38), and derives the explicit link between continuous time and discrete time model parameters. Estimated continuous time parameters implied by the discrete time point estimates in Panel (a) are shown in Panel (b) of Table 7. Intuitively, the main implications obtained for the discrete time estimates carry over to the continuous time parameters. For instance, robust estimation results imply (i) a larger slope coefficient ν of log dividend-price ratios in the (predictive) equation for expected returns, (ii) a more negative correlation parameter ρ and (iii) lower volatilities of both stock returns and log dividend-price ratios.

Given the continuous time-estimates, optimal portfolio policies in Equation (40) can be implemented. Panel (a) of Figure 12 depicts the resulting optimal stock allocations for different time horizons, based on the parameter estimates for the 1929–2009 sample. For each set of estimated parameters, myopic and hedging demands in Panels (b) and (c) are both smaller for the ambiguity averse agent, which is in line with the intuition that standard CRRA investors are more optimistic about future stock performance than ambiguity averse investors. Since the standard CRRA agent is more confident with the predictive power of the dividend yield, she also more rapidly increases the equity allocation, as a function of the investment horizon, due to a larger intertemporal hedging motive.[15]

[Figure 12 about here.]

Robust parameter estimates imply an increased myopic and intertemporal hedging demand, relative to the weights based on standard maximum likelihood estimators. Robust estimates imply both a stronger predictive relation and a more negative correlation between returns and predictor variable, which lead to the larger allocations to equities. Overall, these differences are more pronounced for standard CRRA agents, who do not take the uncertainty into account in selecting their optimal policies.

5.4. Ambiguous Predictability: the Economic Value of Robust Estimation

We conclude our analysis by quantifying the added economic value of our robust estimation approach for dynamic asset allocation under ambiguous predictability features.

5.4.1. Basic Findings

First, consider point estimates of classical and robust estimators from the whole sample. Based on these estimates, we implement the optimal portfolio policy (Equation 40) for an investment horizon of one year, using the current dividend-price ratio and risk-free interest rate in each month t as inputs. Panels (a) and (b) of Table 8 show time series averages of the resulting optimal portfolio weights, for different levels of risk and ambiguity aversion.

[Table 8 about here.]

As expected, the stronger predictability relation and hedging motive implied by robust parameter estimates induces larger average portfolio weights. The economic value of in-sample optimal policies can be evaluated based on realized utility. To this end, the equivalent wealth obtained from investing in the optimal portfolio is computed twelve month after month t. Panels (c) and (d) depict the time series average realized wealth, based on classical and robust parameter estimates, respectively. We find that for every level of γ and ϑ, investing according to policies implied by robust parameter estimates yields a higher wealth than policies based on classical estimates. The gain in utility is largest for low levels of risk aversion (e.g., more than

[15] Assuming an intermediate level of risk aversion in line with Barberis (2000), γ is assumed to be 3. Based on the average T-Bill return, the continuously compounded annual risk-free interest rate equals $r_f = 3.59\%$. Moreover, the initial log dividend-price ratio is set to be the average log dividend-price ratio over the sample period ($\log(3.66\%)$).

2.33% for $\gamma = \vartheta = 2$), but is still economically significant, with at least 0.5% per year, for very risk averse investors.

Policies estimated by robust methods are designed to be less sensitive to particular data features. To shed light on this important aspect, we study the sensitivity of estimated portfolio weights and certainty equivalent wealths along the lines of Dell'Aquila, Ronchetti, and Trojani (2003), who perturb the most influential observations, identified from estimated Huber weights, to investigate the robustness of the relevant decision variable. As a simple example, consider a perturbation of the October 1987 return, the most extreme postwar observation in the sample, by up to ±30%. This corresponds to varying that month's actual return of -28.52% between -37.08% and -19.96%. The top panel of Table 9 shows the resulting minimum and maximum of the time series means of portfolio weights, after perturbing the three most influential returns by ±30%.

[Table 9 about here.]

We find that for both standard ($\vartheta = 0$) and robust ($\vartheta = 3$) agents, portfolio weights based on classical estimates are much more sensitive to such a perturbation of only approximately 0.3% of all observations. While robust estimates only imply a change in average weights of maximally 1.34% (0.90% for $\vartheta = 3$), classical parameter estimates induce a difference between weights of 3.71%, as a result of the ±30% perturbations (2.13% for $\vartheta = 3$): The sensitivity of optimal portfolio weights to this small perturbation is almost three times as large for non-robust parameter estimates. These findings are confirmed, and in some cases stronger, for perturbations of ±10% or perturbations of the six most influential observations; see the remaining panels of Table 9. Large portfolio weight sensitivities are directly linked to large sensitivities of the corresponding optimal utility or certainty equivalent wealth. Table 9 confirms this intuition: The average wealth is more sensitive to a limited fraction of observations when agents invest based on classical parameter estimates. The largest cross sectional differences arise when perturbing the six most influential returns: The minimum average wealth of the utility maximizer using classical estimates is 1.21% smaller than the maximum. This difference is less than half as big when the agent estimates her reference model with robust methods. Moreover, the minimum average realized wealth using robust parameters is always larger than the maximum average realized wealth obtained by relying on classical estimates. This feature holds true regardless of the size of the perturbation and the level of ambiguity aversion. Thus, our finding that using robust estimators to recover the predictability structure leads to a larger realized wealth is confirmed, even when the most influential returns are perturbed.

5.4.2. A Simple Real-Time Exercise

In reality, the agent needs to estimate model parameters in real time, in order to derive her investment policy. Thus, in a more realistic setup she will add one observation to her sample in each month t and estimate the reference model, leading to an additional potential variation in the portfolio allocation. On the one hand, each new observation can contain valuable information about economic fundamentals and potential return predictability features. On the other hand, a newly incorporated return observation may simply come from an unknown time-varying feature of the underlying data generating process and be very influential for classical maximum likelihood estimators. In this case, the implied portfolio weights can actually result in a lower out-of-sample utility level.

Table 10 depicts the average portfolio weights and their standard deviations, implied by classical and ambiguity averse utility maximizers with robust parameter estimates, in an increasing sample starting in 1929 with initially 300 observations. Similarly, the one year ahead realized wealth from following the portfolio strategies out-of-sample is shown for different levels of risk aversion γ and degree of ambiguity aversion ϑ.

[Table 10 about here.]

In the top panel, the agent's reference model incorporates return predictability. We find that the average realized out-of-sample wealth implied by robust parameter estimates dominates the one of classical parameter estimation methods, regardless of the degree of risk and ambiguity aversion. The realized differences in wealth can be economically relevant, with differences of about 0.87% per year in some cases. For comparison, the second panel from the top assumes instead a random walk process for returns. Here, the increase in wealth from using robust estimators is even larger: For instance, the classic utility maximizer with $\gamma = 5$ can improve her realized wealth by more than 1.3% per year, when she estimates her reference model with a robust approach. The last four panels from the top consider mixed strategies, in which at each month t the agent first tests the predictability hypothesis, at a given confidence level between 5% and 30%. When she rejects the null of no predictability, she applies the corresponding dynamic optimal policies, in the attempt to exploit predictability features. In all other cases, she applies the myopic policy implied by a random walk assumption. The results show that, when using robust estimators and confidence intervals, an investor with a good degree of confidence in the predictability hypothesis, e.g., above 10%, outperforms strategies that dogmatically believe either in the predictability or the random walk hypothesis. The resulting differences in out-of-sample wealth can be economically relevant in some cases. For instance, an expected utility maximizer with a confidence of 20% in the predictability hypothesis and risk aversion $\gamma = 5$ can improve her wealth by about 0.5% per year, relative to a simple random walk assumption. On the other hand, mixed strategies implied by classical estimators do not provide economically significant improvements. Table 11 provides an explanation for these results, by presenting average portfolio weights and realized wealth for different subsamples.

[Table 11 about here.]

Recalling the rolling regression evidence for predictability in Figure 10, we find that dynamic policies including predictability dominate policies implied by the random walk assumption in the middle period (1973–1990), when the evidence for predictability is highest: In that case, an agent who believes in predictability outperforms the random walk investor. However, in the most recent decades and the earliest subsample, the evidence for predictability is weaker and the random walk policy dominates the dynamic one. It follows, that the mixed strategy implied by our robust estimation approach is better able to exploit predictability features in the data, precisely when the agent can be more confident on (ambiguous) predictability features based on real-time data information.

Overall, the findings from this simple real-time portfolio allocation exercise corroborate the in-sample results: For all levels of risk and ambiguity aversion, different choices of reference model and almost all samples, policies implied by robust estimates lead to greater wealth.

6. Conclusion

Small unknown time-varying features in dynamic financial settings can be difficult to identify statistically and to model theoretically, motivating aversion to ambiguity as a convenient assumption to describe economic agents' behavior. The necessity to estimate models in presence of unknown time-varying features further complicates the task of economic agents who want to develop robust decision rules and model builders, trying to quantify key equilibrium variables like, e.g., risk and ambiguity premiums. In these contexts, we introduce a widely applicable robust estimation approach, which is characterized by a bounded sensitivity of estimated optimal policies and general equilibrium parameters to unknown time-varying features in the underlying data generating process. We find that such a bounded influence estimation approach is key for producing (i) estimated optimal policies that are robust to unknown time-varying features and (ii) estimated equilibrium variables that are more consistent with the assumption of ambiguity aversion in general equilibrium. First, in consumption and portfolio planning problems with ambiguity aversion, small unknown time-varying features, in the conditional mean of returns or the probability structure of rare events, can generate economically relevant utility losses, which are successfully bounded by our robust estimation approach. Second, within general equilibrium economies, unknown time-varying rare event features can severely bias estimates of risk or ambiguity premiums, produced by standard estimation approaches from cross sectional and time series information on underlying's returns and derivative prices. These biases are virtually eliminated by our bounded influence estimation approach. Finally, in a real data study on portfolio choice with ambiguous predictability, our approach uncovers predictability structures that consistently produce both (i) a larger out-of-sample utility than classical approaches and (ii) optimal portfolio weights more robust to abnormal data structures. Moreover, when focusing on real-time portfolio strategies estimated by our robust method, a mixed strategy of an agent with a realistic degree of confidence in the predictability hypothesis can produce larger utilities than the strategy implied by a random walk assumption for market returns.

Overall, these findings and the wide applicability of our robust approach to, e.g., (pseudo) maximum likelihood, generalized method of moments and efficient method of moments settings, suggest the usefulness of our methodology more generally, to estimate robust optimal policies and general equilibrium parameters in a broad variety of dynamic settings of ambiguity aversion. Such applications can potentially produce a number of new insights and interpretations for the growing literature on ambiguity aversion in finance, which largely abstracts from the implications of robust estimation in studying the consequences of ambiguity aversion for asset pricing.

Appendix A. Updating Algorithm for Robust Auxiliary Estimator

The robust auxiliary estimator can be computed using the following algorithm: In a first step, the initial estimate $\boldsymbol{\theta}^{(0)}$ obtained from classical PML estimation is used to compute $\mathbf{A}^{(0)}$:

$$\mathbf{A}^{(0)\top}\mathbf{A}^{(0)} = \left[\frac{1}{T}\sum_{t=1}^{T}\mathbf{s}\left(\mathbf{y}_{t-m+1}^{t},\boldsymbol{\theta}^{(0)}\right)\mathbf{s}\left(\mathbf{y}_{t-m+1}^{t},\boldsymbol{\theta}^{(0)}\right)^{\top}\right]^{-1}. \tag{A.1}$$

Next, the matrix $\mathbf{A}$ is updated by including the weights of Equation (17) based on matrix $\mathbf{A}^{(0)}$ and:

$$\begin{aligned}\mathbf{A}^{(1)\top}\mathbf{A}^{(1)} = \Bigg[&\frac{1}{T}\sum_{t=1}^{T}\mathbf{s}\left(\mathbf{y}_{t-m+1}^{t},\boldsymbol{\theta}^{(0)}\right)\mathbf{s}\left(\mathbf{y}_{t-m+1}^{t},\boldsymbol{\theta}^{(0)}\right)^{\top} \\ &\times\left(\min\left(\mathbf{1}, c\left\|\mathbf{A}^{(0)}\mathbf{s}(\mathbf{y}_{t-m+1}^{t},\boldsymbol{\theta}^{(0)})\right\|^{-1}\right)\right)^{2}\Bigg]^{-1}.\end{aligned} \tag{A.2}$$

Subsequently, robust parameter estimates based on matrix $\mathbf{A}^{(1)}$ can be computed by solving:

$$\frac{1}{T}\sum_{t=1}^{T}\mathbf{s}\left(\mathbf{y}_{t-m+1}^{t},\boldsymbol{\theta}^{(1)}\right)\times\min\left(\mathbf{1}, c\left\|\mathbf{A}^{(1)}\mathbf{s}(\mathbf{y}_{t-m+1}^{t},\boldsymbol{\theta}^{(0)})\right\|^{-1}\right) = \mathbf{0}. \tag{A.3}$$

Finally, the last two steps need to be iterated until convergence of the robust parameter estimates.

Appendix B. Reference Model Estimation

Appendix B.1. Reference Model

The agent considers the reference model:

$$dS_t = \mu_S S_t dt + \sigma S_t dB_t + \left(e^{\xi^Y} - 1\right) S_t dN_t, \tag{B.4}$$

where $\mu_S = \mu + \frac{1}{2}\sigma^2$, and ξ^Y is the jump size corresponding to the Poisson process with intensity λ_Y; $\xi^Y \sim N(\mu_Y, \sigma_Y^2)$.

Appendix B.2. Estimation Methodology

Letting $f_p(\cdot)$ and $\phi(\cdot)$ denote the Poisson and Normal probability density, respectively, Model (B.4) can be estimated by maximizing the following log-likelihood function:

$$\begin{aligned}
l(\boldsymbol{\theta}|\mathbf{y}) = \ln L(\boldsymbol{\theta}|\mathbf{y}) &= \ln\left(\prod_{t=1}^{T}\left(\sum_{n=0}^{\infty} f_p\left(n|\lambda_Y\right)\cdot\phi\left(y_t|\bar{y},\sigma_{t,tot}^2\right)\right)\right) && \text{(B.5)}\\
&= \sum_{t=1}^{T}\left(\ln\left(\sum_{n=0}^{\infty} f_p\left(n|\lambda_Y\right)\cdot\phi\left(y_t|\bar{y},\sigma_{t,tot}^2\right)\right)\right) && \text{(B.6)}\\
&= \sum_{t=1}^{T}\ln\left(L_t(\boldsymbol{\theta}|y_t)\right) = \sum_{t=1}^{T} l_t(\boldsymbol{\theta}|y_t), && \text{(B.7)}
\end{aligned}$$

where

$$\begin{aligned}
\sigma_{t,tot}^2 &= (\sigma^2 + n\sigma_Y^2), && \text{(B.8)}\\
\bar{y} &= (\mu + n\mu_Y). && \text{(B.9)}
\end{aligned}$$

Appendix C. Figures and Tables

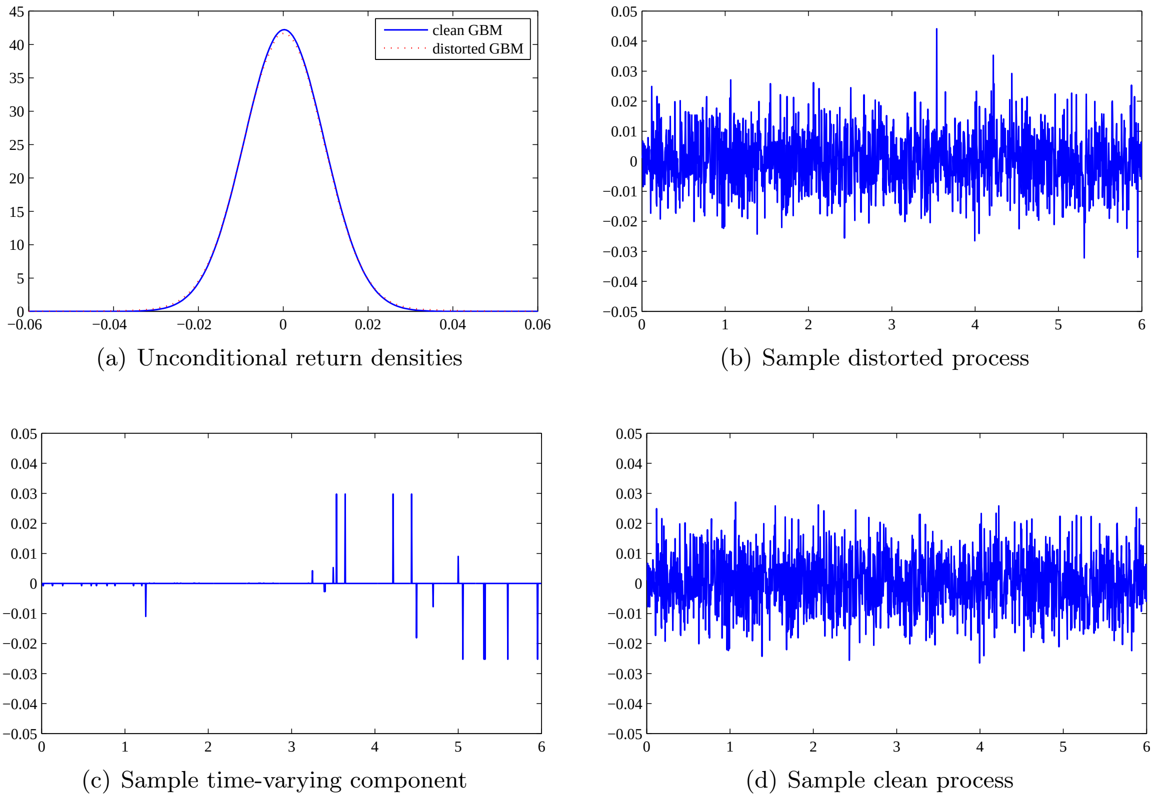

(a) Unconditional return densities

(b) Sample distorted process

(c) Sample time-varying component

(d) Sample clean process

Figure 1: Panel (a) depicts densities implied by the approximate GBM of Equation (5) and the corresponding GBM. The figure shows the return densities of the true DGP and the agent's reference model based on 1,000,000 observations generated from the stock process specified in Equation (5); $\mu = 0.06$, $\sigma = 0.15$. A sample realization (six years) of the approximate GBM is shown in Panel (b). The components of this process, a time-varying drift distortion and a clean GBM, are shown in Panels (c) and (d), respectively.

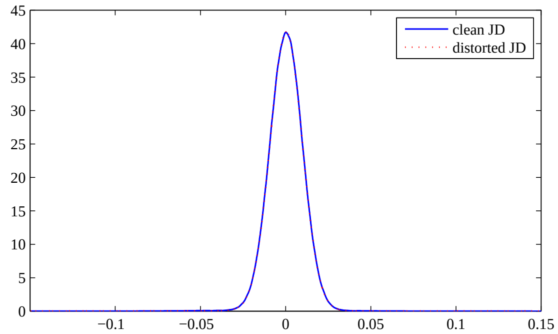

(a) Unconditional return densities

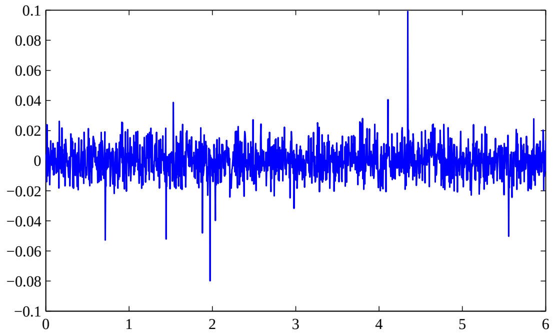

(b) Sample distorted process

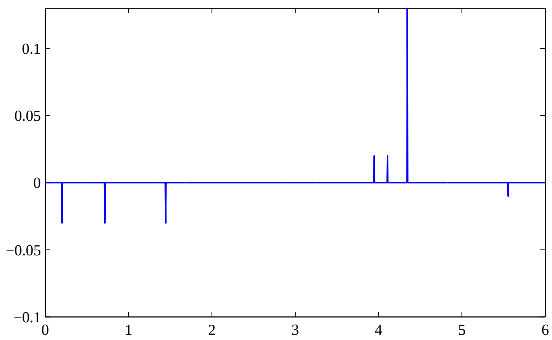

(c) Sample time-varying component

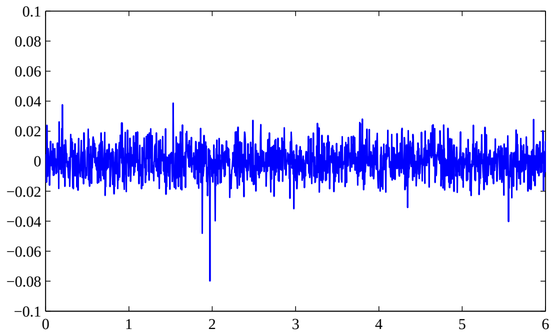

(d) Sample clean process

Figure 2: Panel (a) depicts densities implied by the approximate JD of Equation (25) and the corresponding JD model. The Figure shows the return densities of the true DGP and the agent's reference model based on 1,000,000 observations generated from the stock process specified in (25); $\mu = 0.08$, $\sigma = 0.15$, $\lambda_Y = 3$, $\mu_Y = -0.01$, $\sigma_Y = 0.04$. A sample realization (six years) of the approximate JD model is shown in Panel (b). The components of this process, a time-varying jump size distortion and a clean JD model, are shown in Panels (c) and (d), respectively.

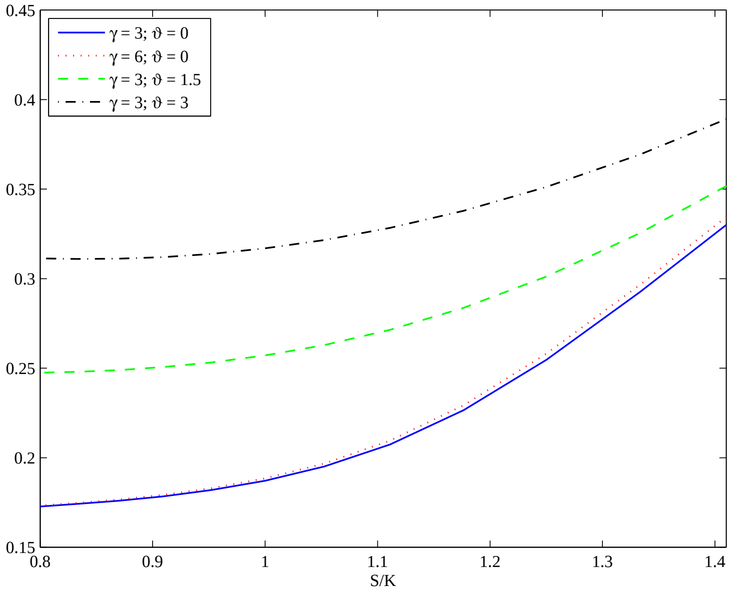

Figure 3: Equilibrium option prices (quoted in terms of Black-Scholes implied volatilities) for different levels of risk and ambiguity aversion of the representative agent. For each scenario, one-year European call option prices are computed using the equilibrium model for varying degrees of moneyness. The parameters of the jump-diffusion are $\mu = 0.08$, $\sigma = 0.15$, $\lambda_Y = 3$, $\mu_Y = -0.01$, $\sigma_Y = 0.04$.

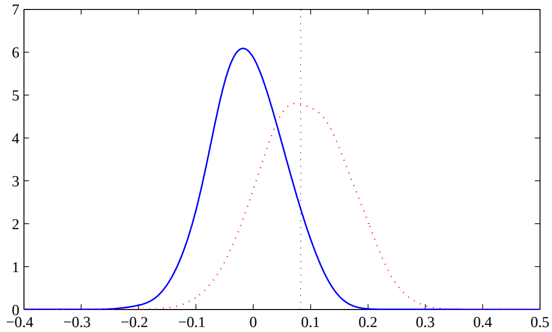

(a) Total equity premium; distorted data

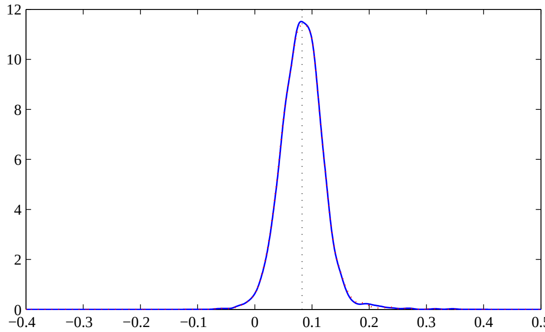

(b) Total equity premium; clean data

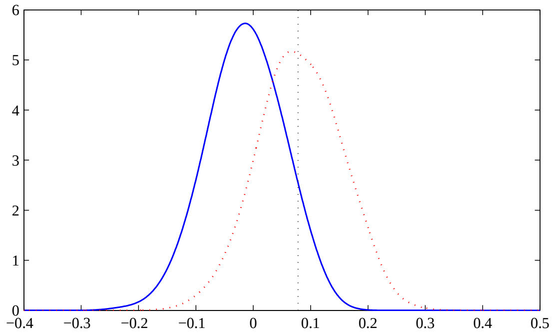

(c) Risk premium; distorted data

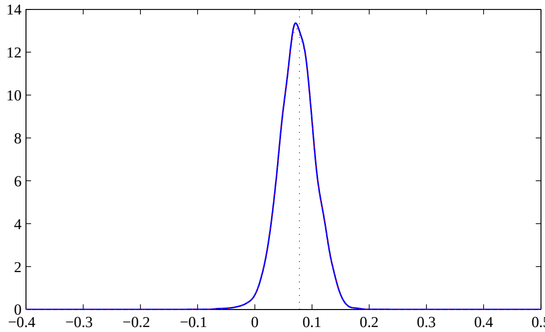

(d) Risk premium; clean data

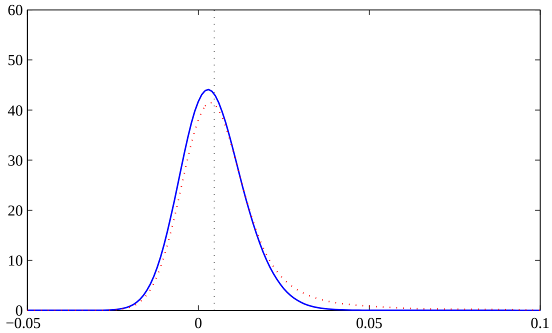

(e) Ambiguity premium; distorted data

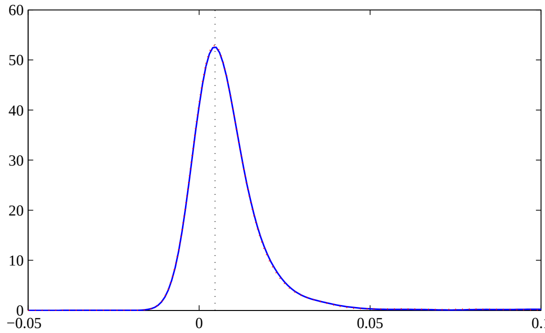

(f) Ambiguity premium; clean data

Figure 4: Equity premium for distorted and clean data according to Equations (27) and (29). Panels (a) and (b) show the total equity premium of an economy with ambiguity aversion. The vertical line indicates the true premium ignoring parameter estimation. The blue, solid lines show estimated premiums based on classical parameter estimates, and the red, dashed lines give the distributions corresponding to robust estimates. Panels (c)–(f) decompose the total premium into risk and ambiguity components: Panels (c) and (d) show the risk premiums corresponding to a jump-diffusion economy and Panels (e) and (f) report the difference between total premiums and the risk premiums. The three plots on the left are obtained from distorted data in accordance with Equation (20) and equilibrium option prices computed from the partial differential equation in (30). The true preference parameters are $\gamma = 2$ and $\vartheta = 1$; the jump-diffusion is characterized by $\mu = 0.08$, $\sigma = 0.15$, $\lambda_Y = 3$, $\mu_Y = -0.01$, $\sigma_Y = 0.04$.

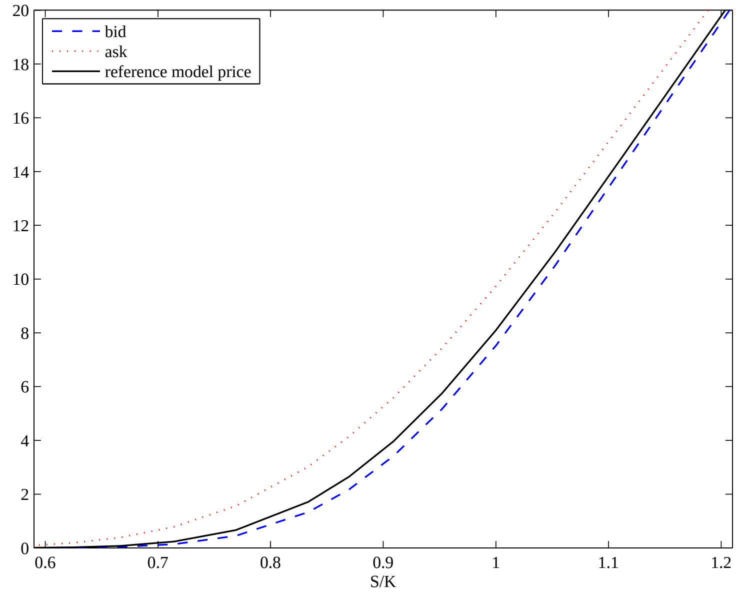

(a) Bid and ask prices as a function of moneyness

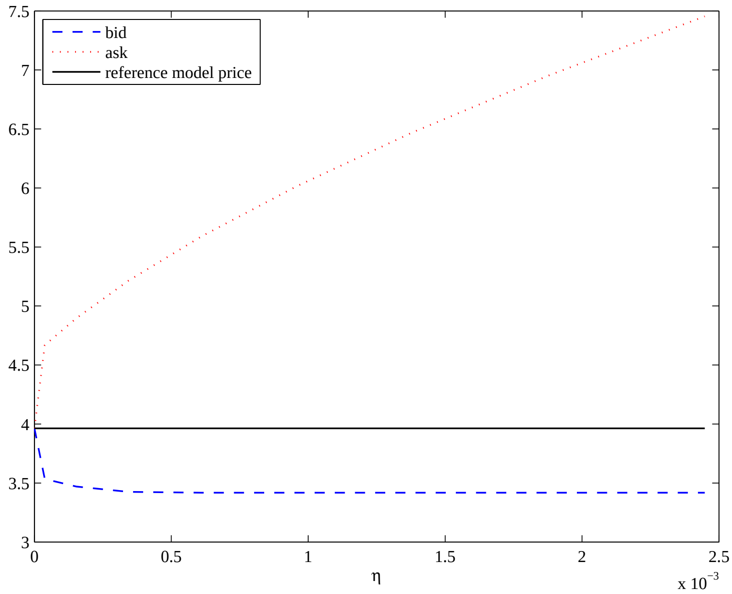

(b) Bid and ask prices as a function of ambiguity

Figure 5: Panel (a) shows bid and ask prices as a function of moneyness. The entropy bound η in Equations (31) and (32) equals 0.000612. The time to maturity of the option is 1 year, $S(0) = 100$, and $r = 3\%$. In Panel (b) bid and ask prices for different levels of ambiguity are depicted. The strike price equals $K = 110$. For comparison, the solid line in both panels correspond to the price of the option without ambiguity.

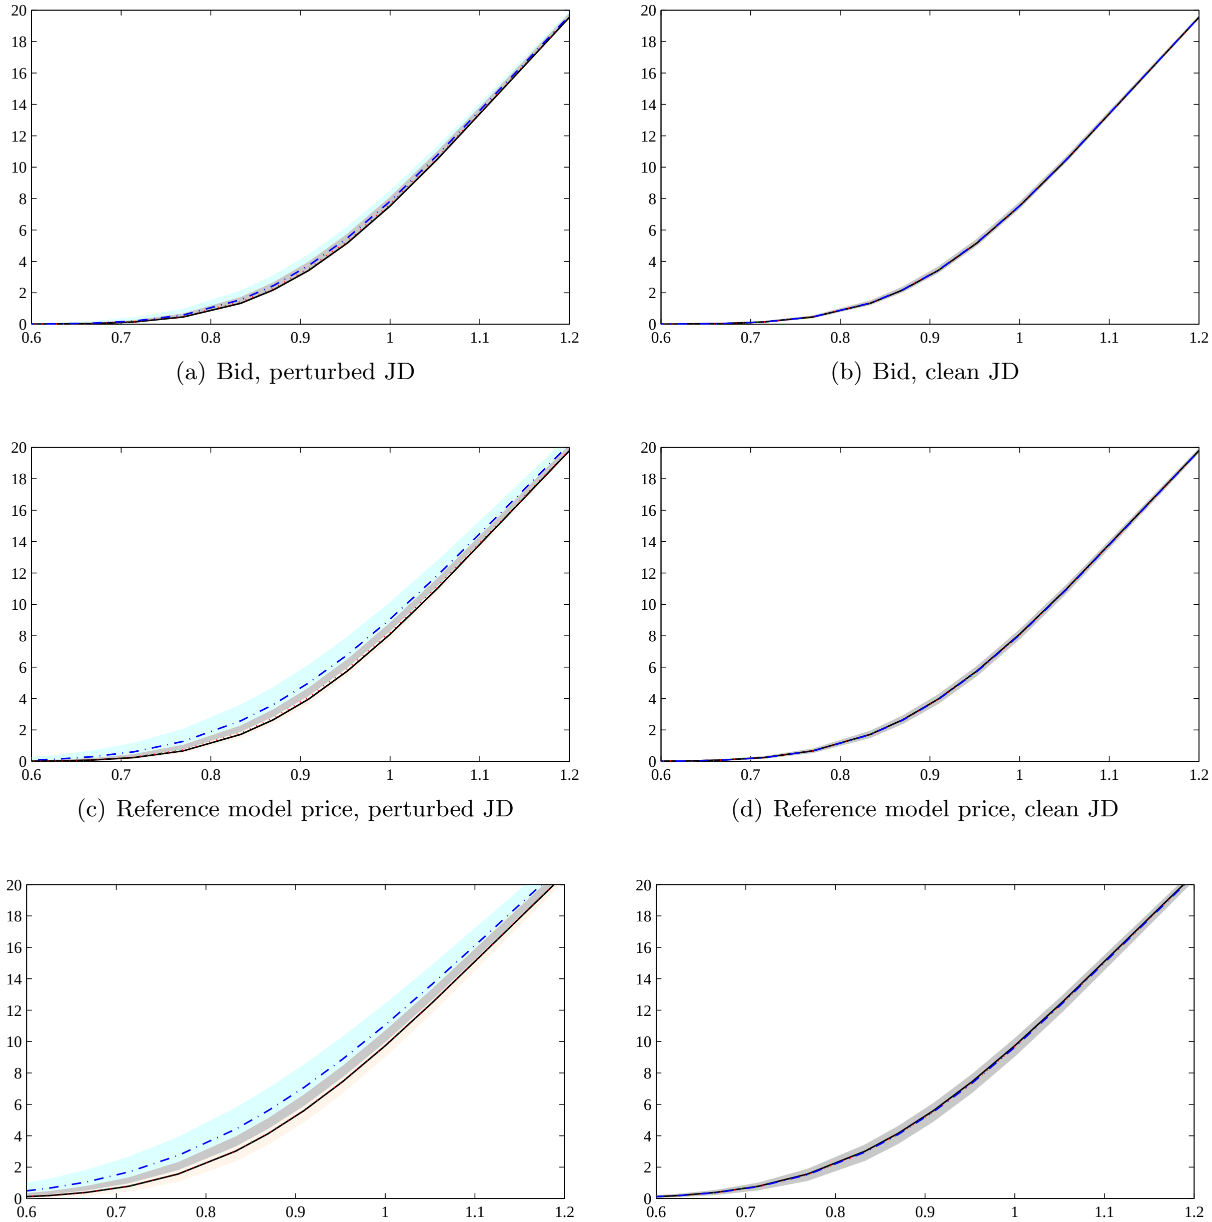

(a) Bid, perturbed JD

(b) Bid, clean JD

(c) Reference model price, perturbed JD

(d) Reference model price, clean JD

(e) Ask, perturbed JD

(f) Ask, clean JD

Figure 6: Bid and ask prices as a function of moneyness for estimated parameters. In Panels (a), (c), and (e), estimates are obtained based on data from a perturbed jump diffusion (Equation (25)). Panels (b), (d), and (f) show bid and ask quotes for estimates from a clean jump diffusion. The red, dotted and the blue, dashed lines denote the robust and classical median estimate, respectively. The shaded areas mark the 5% and 95% quantiles. For comparison, the black, solid line in all panels corresponds to the price of the option based on true parameters. The time to maturity of the option is 1 year, $S(0) = 100$, and $r = 3\%$. The entropy bound is $\eta = 0.00061$.

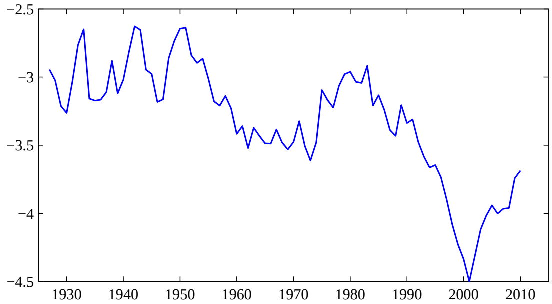

(a) Dividend-price ratio

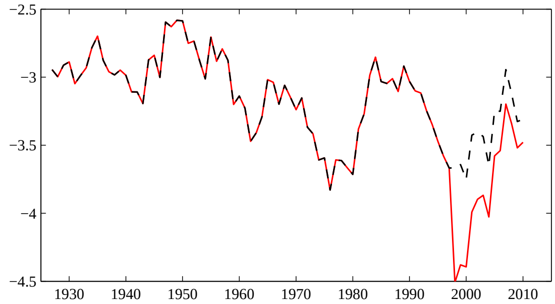

(b) Simulated dividend-price series

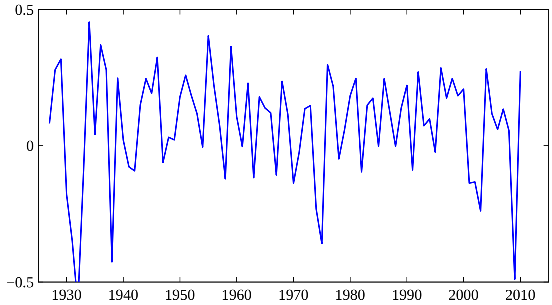

(c) Returns

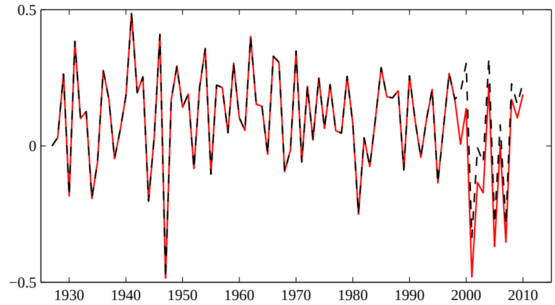

(d) Simulated returns

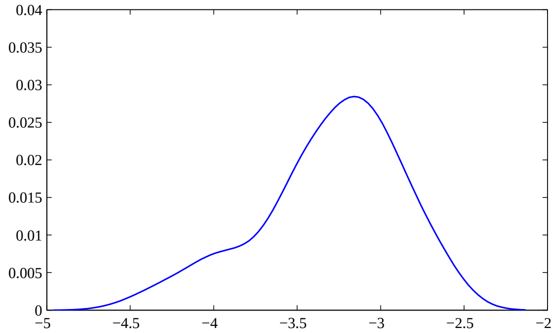

(e) Empirical distribution function of dividend-price ratio

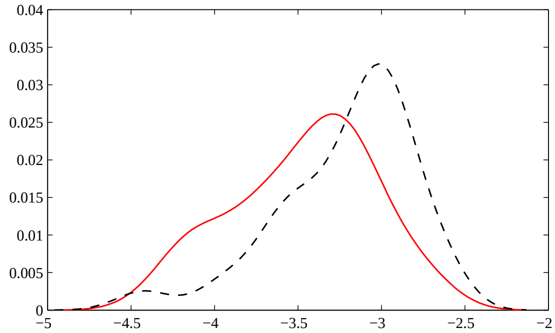

(f) Empirical distribution function of simulated dividend-price ratios

Figure 7: Panels (a) and (c) show the dividend-price ratios and excess returns of the CRSP value-weighted market portfolio (including NYSE, AMEX and NASDAQ) for the period 1926–2009. Simulated sample paths for the dividend-price ratios and excess returns are shown in Panels (b) and (d), respectively. These paths are obtained by simulating returns and log dividend-price ratios according to Equations (34) and (34) with parameters estimated from observed data during 1927–1994. The dashed line in panel (b) and (d) corresponds to sample paths generated exactly according to Equations (34) and (34). The simulation underlying the solid line is almost identical, except that $\omega = -0.40$ is replaced by $\tilde{\omega} = -1.25$ at a single point in time (in 1997). The empirical distribution functions for the observed log dividend-price ratio as well as the simulated series are plotted in Panels (e) and (f).

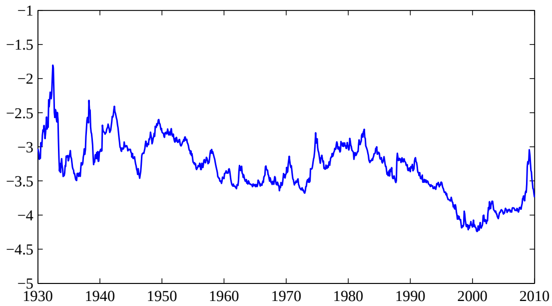

(a) Monthly dividend-price ratio

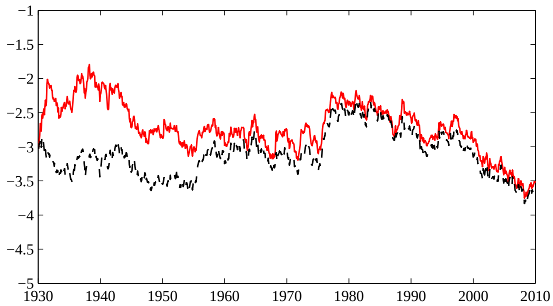

(b) Simulated monthly dividend-price series

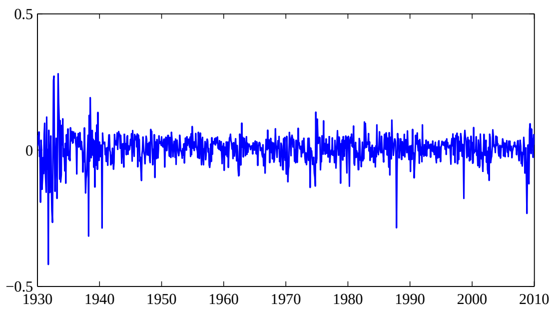

(c) Monthly returns

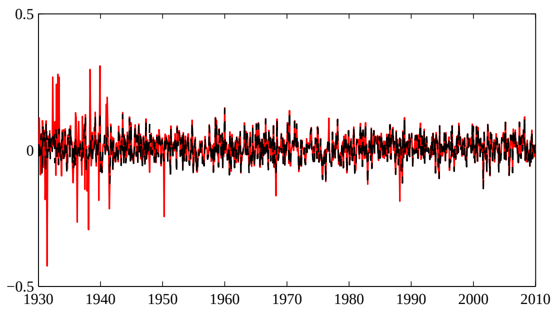

(d) Simulated monthly returns

Figure 8: Panels (a) and (c) show monthly dividend-price ratios and excess returns of the CRSP value-weighted market portfolio (including NYSE, AMEX and NASDAQ) for the period 12/1929–12/2009. Simulated sample paths for the dividend-price ratios and excess returns are shown in Panels (b) and (d), respectively. These paths are obtained by simulating returns and log dividend-price ratios according to Equations (34) and (34). The dashed line in panel (b) and (d) corresponds to sample paths generated exactly according to Equations (34) and (34). The simulation underlying the solid line is almost identical, except that jumps are added randomly to the return and price-dividend series. In the first 15% of the simulated observations, jumps are added to the process with an intensity of four jumps per year, mean jump size equal to zero, and a jump variance of 0.15^2. In the remaining 85% of the observations, jumps occur only once every 8 years, but tend to be more negative (mean jump size -0.05, jump variance 0.15^2).

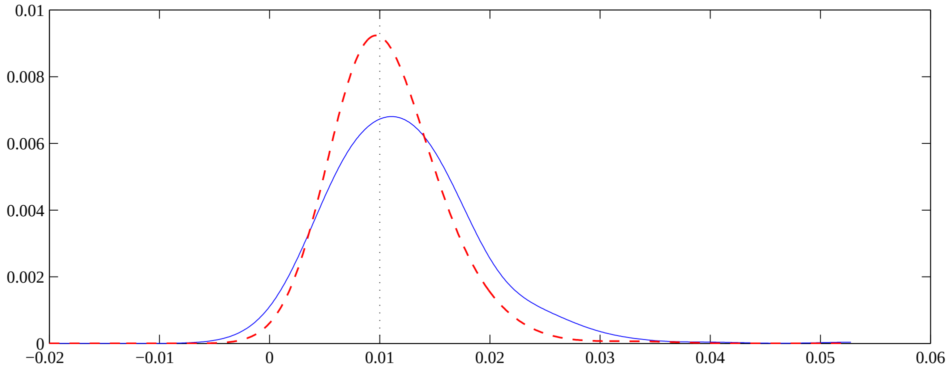

(a) Density of point estimates for predictability parameter (β)

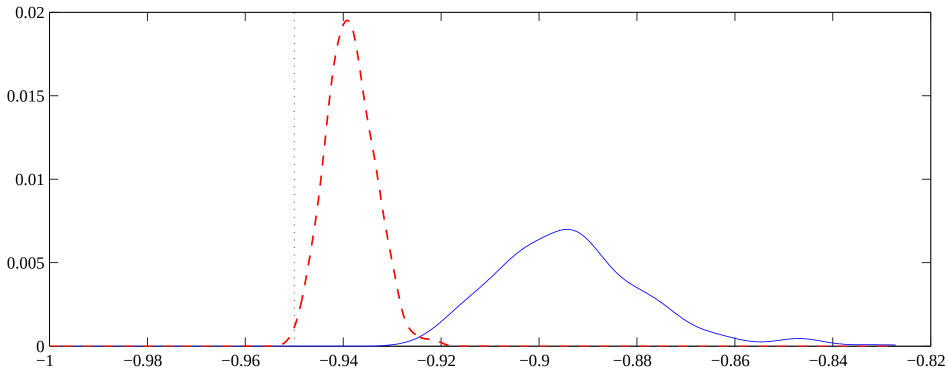

(b) Density of point estimates for correlation parameter (ϱ)

Figure 9: Panels (a) and (b) show the empirical distributions of point estimates for the predictability coefficient β and the correlation of the innovation ϱ from Equations (34) and (34), respectively. The solid lines correspond to classical maximum likelihood estimates, whereas the dashed lines depict the EDFs of robust parameter estimates ($c_{Huber} = 5$). The estimates are based on 10,000 simulated sample paths (1,000 observations) which are obtained by simulating returns and log dividend-price ratios according to Equations (34) and (34). In the first 15% of the simulated observations, jumps are added to the process with an intensity of four jumps per year, mean jump size equal to zero, and a jump variance of 0.15^2. In the remaining 85% of the observations, jumps occur only once every 8 years, but tend to be more negative (mean jump size -0.05, jump variance 0.15^2).

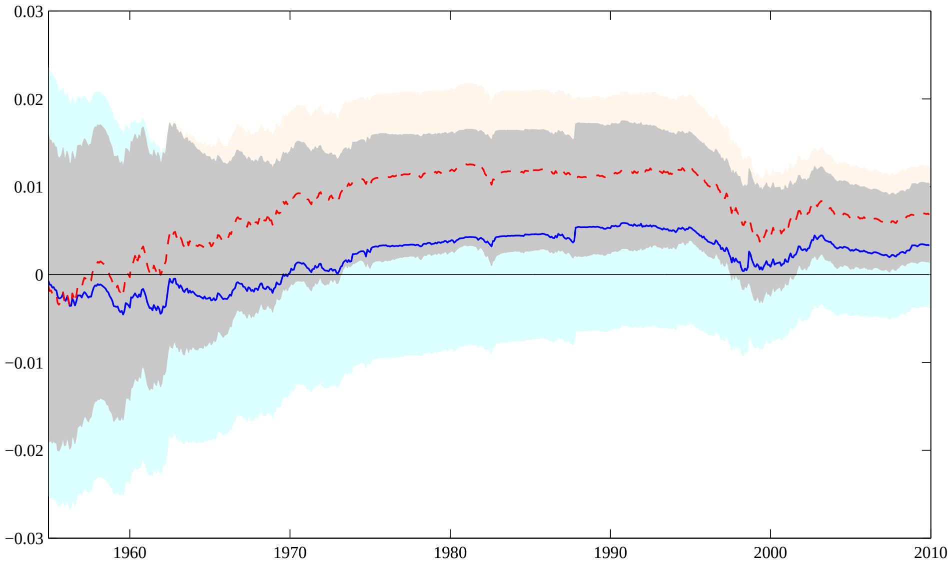

Figure 10: Evolution of d_3 point estimates together with 68% convidence intervals ($\pm$ one standard error) over time. In every month t, Model (35) is estimated based on sample data from January 1929 up to the current month. The red, dotted line corresponds to robust point estimates ($c_{Huber} = 6$); the union of the top and middle shaded area depicts the robust confidence intervals. The solid, blue line represents the classical point estimates; the classical confidence interval is given by the union of the bottom and middle shaded areas.

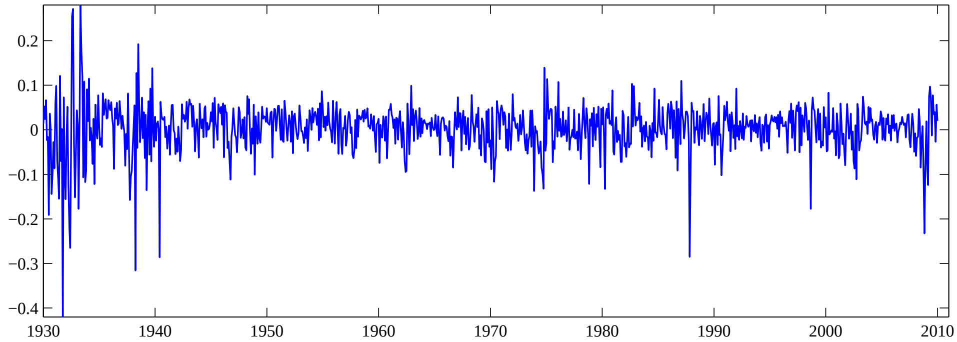

(a) Excess stock returns

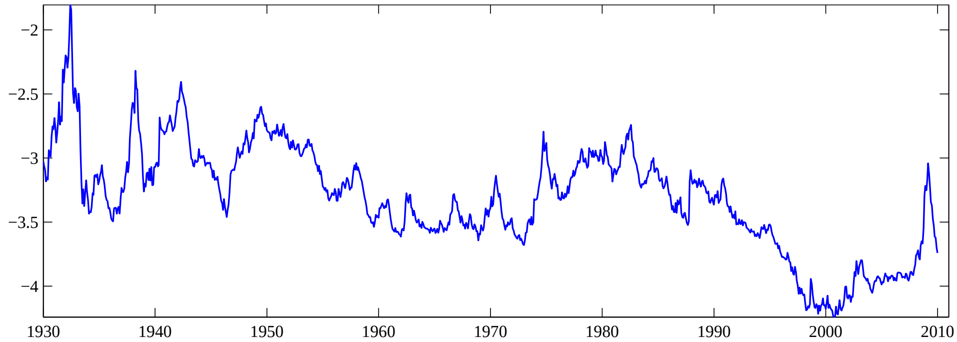

(b) Dividend-price ratio

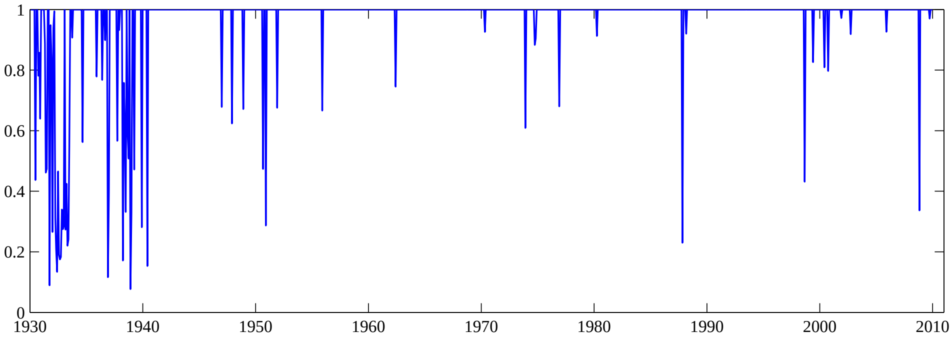

(c) Huber weights

Figure 11: Panels (a) and (b) show excess stock returns and the dividend-price ratio of the CRSP value-weighted market portfolio (including NYSE, AMEX and NASDAQ) for the period 1929–2009, respectively. The Huber weights corresponding to the observations are depicted in Panel (c). The econometric constant c_{Huber} used in the estimation equals 6.

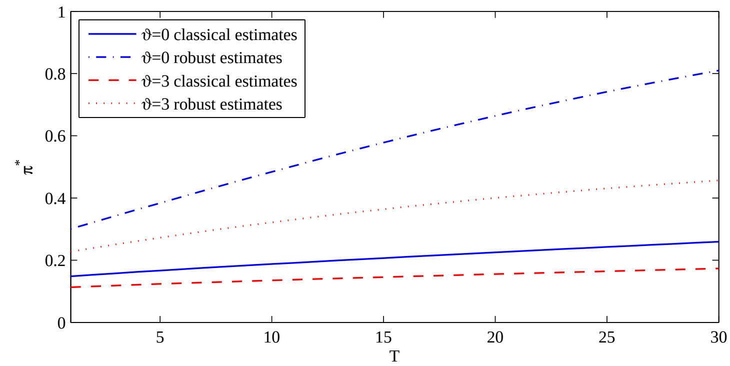

(a) Total demand

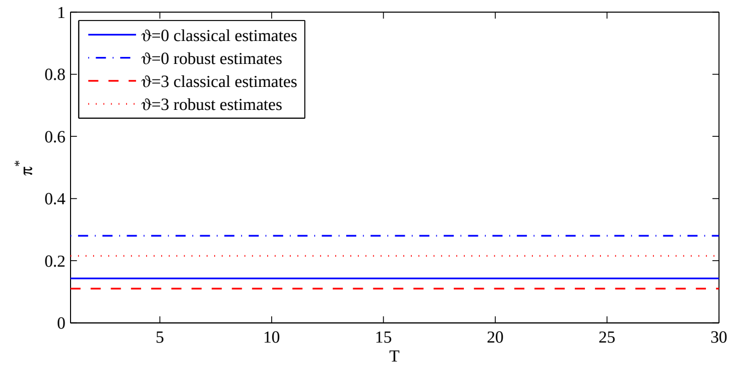

(b) Myopic demand

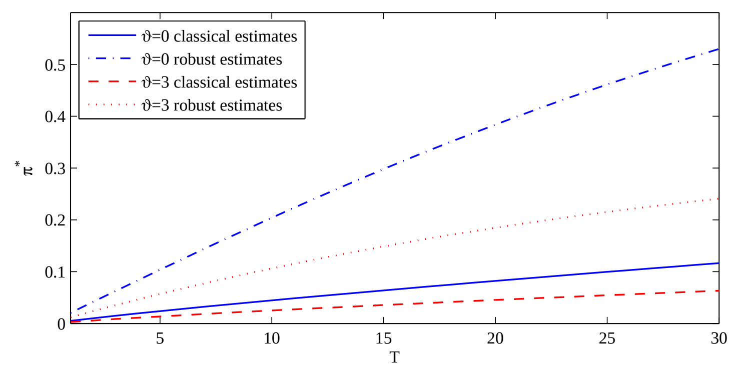

(c) Hedging demand

Figure 12: Portfolio allocation as function of investment horizon. Panel (a) shows the optimal allocation into equities π^* of classical and uncertainty averse investors as function of investment horizon for both classical and robust parameter estimates. The myopic and hedging demand components of total demand are shown in Panels (b) and (c), respectively. Continuously compounded annual risk-free interest rate $r_f = 3.59\%$, coefficient of risk aversion $\gamma = 3$, initial log dividend-price ratio equals $\log(3.66\%)$.

Table 1: Expected and Realized Utility

	Ex-ante: Expectation		*Ex-post:* True realized utility			
	(1)	(2)	(3)	(4)	(5)	(6)
In-sample	GBM	distorted	GBM	GBM	distorted	distorted
Out-of-sample	*	*	GBM	distorted	GBM	distorted
			Panel (a): True parameters			
Exp. utility maximizer	1,047	1,047	1,047	926	1,047	926
Robust agent						
$\vartheta = 0.2$	919	919	1,047	934	1,047	934
$\vartheta = 0.6$	801	801	1,046	947	1,046	947
$\vartheta = 1.0$	747	747	1,046	958	1,046	958
$\vartheta = 2.0$	699	699	1,046	976	1,046	976
$\vartheta = 3.0$	690	690	1,046	988	1,046	988
			Panel (b): (P)MLE			
Exp. utility maximizer	973	926	985	601	947	476
Robust agent						
$\vartheta = 0.2$	619	491	999	641	969	515
$\vartheta = 0.6$	485	376	1,015	702	998	594
$\vartheta = 1.0$	447	348	1,024	763	1,019	665
$\vartheta = 2.0$	422	334	1,034	882	1,034	798
$\vartheta = 3.0$	421	335	1,038	942	1,040	872
			Panel (c): Robust estimates			
Exp. utility maximizer	967	949	980	586	961	538
Robust agent						
$\vartheta = 0.2$	606	550	991	627	978	576
$\vartheta = 0.6$	474	425	1,009	694	1,003	655
$\vartheta = 1.0$	437	394	1,022	757	1,021	725
$\vartheta = 2.0$	414	376	1,033	881	1,034	838
$\vartheta = 3.0$	413	377	1,037	940	1,039	908

Notes: The table shows the *ex-ante* (worst case) utility expectations and the simulation based average *ex-post* realized utilities in wealth equivalents of an expected utility maximizer and her robust counterpart. When knowing the parameters of the reference model, Columns (1) & (2) of Panel (a) give the *ex-ante* expectations of an agent who considers a geometric Brownian motion (GBM) as reference model. The average *ex-post* wealth equivalents, given the implemented policies, are shown in Columns (3) to (6). Column (3) reports the wealth equivalents for the case that a GBM generates the data, both in- and out-of-sample. The numbers in Column (4) are based on an out-of-sample DGP that is a symmetric perturbation to the GBM's drift (Equation (6)). Column (5) further considers in-sample distortion and a GBM out-of-sample. In Column (6) data have been generated from two different types of distortions in- and out-of-sample, respectively. Panels (b) & (c) show the same six cases but recognize the need to estimate model parameters from the in-sample data. For this purpose, Panel (b) reports wealth equivalents based on maximum likelihood estimates and Panel (c) considers the robust estimator. Portfolio weights for realized and expected utilities are computed based on five years of daily data generated by the stock process specified in Equations (5) and (6); $\mu = 0.06$, $\sigma = 0.15$. The econometric constant c_{Huber} used in the estimation equals 1.345 to achieve 95% efficiency if the data actually have been generated by the GBM reference model. The number of simulations equals 10,000.

Table 2: Expected and realized utility for distortion with opposite sign

	Ex-ante: Expectation	*Ex-post:* True realized utility		
	(1)	(2)	(3)	(4)
In-sample	distorted	GBM	distorted	distorted
Out-of-sample	*	distorted	GBM	distorted
Panel (a): True parameters				
Exp. utility maximizer	1,047	1,156	1,047	1,156
Robust agent				
$\vartheta = 0.2$	919	1,150	1,047	1,150
$\vartheta = 0.6$	801	1,139	1,046	1,139
$\vartheta = 1.0$	747	1,130	1,046	1,130
$\vartheta = 2.0$	699	1,113	1,046	1,113
$\vartheta = 3.0$	690	1,102	1,046	1,102
Panel (b): (P)MLE				
Exp. utility maximizer	978	971	983	782
Robust agent				
$\vartheta = 0.2$	676	994	995	822
$\vartheta = 0.6$	540	1,031	1,009	890
$\vartheta = 1.0$	499	1,055	1,018	930
$\vartheta = 2.0$	471	1,075	1,029	982
$\vartheta = 3.0$	470	1,078	1,033	1,003
Panel (c): Robust estimates				
Exp. utility maximizer	976	972	975	877
Robust agent				
$\vartheta = 0.2$	649	991	989	914
$\vartheta = 0.6$	513	1,025	1,010	960
$\vartheta = 1.0$	474	1,048	1,022	992
$\vartheta = 2.0$	447	1,065	1,032	1,021
$\vartheta = 3.0$	446	1,074	1,036	1,043

Notes: The table shows the *ex-ante* (worst case) utility expectations and the simulation based average *ex-post* realized utilities in wealth equivalents of an expected utility maximizer and her robust counterpart. When knowing the parameters of the reference model, Columns (1) & (2) of Panel (a) give the *ex-ante* expectations of an agent who considers a geometric Brownian motion (GBM) as reference model. The average *ex-post* wealth equivalents, given the implemented policies, are shown in Columns (3) to (6). Column (3) reports the wealth equivalents for the case that a GBM generates the data, both in- and out-of-sample. The numbers in Column (4) are based on an out-of-sample DGP that is a symmetric perturbation to the GBM's drift (Equation (6)). Column (5) further considers in-sample distortion and a GBM out-of-sample. In Column (6) data have been generated from two different types of distortions in- and out-of-sample, respectively. Panels (b) & (c) show the same six cases but recognize the need to estimate model parameters from the in-sample data. For this purpose, Panel (b) reports wealth equivalents based on maximum likelihood estimates and Panel (c) considers the robust estimator. Portfolio weights for realized and expected utilities are computed based on five years of daily data generated by the stock process specified in Equation(13) and $\alpha_t' = -\alpha_t$ from Equation (6);; $\mu = 0.06$, $\sigma = 0.15$. The econometric constant c_{Huber} used in the estimation equals 1.345 to achieve 95% efficiency if the data actually have been generated by the GBM reference model. The number of simulations equals 10,000.

Table 3: Realized utility in the jump-diffusion case

	Degree of robustness					
	$\vartheta = 0$	$\vartheta = 0.2$	$\vartheta = 0.6$	$\vartheta = 1.0$	$\vartheta = 2.0$	$\vartheta = 3.0$
Panel (a): In- and out-of-sample JD						
True parameters	1,056	1,056	1,055	1,055	1,054	1,054
(P)MLE	1,034	1,037	1,040	1,042	1,044	1,046
Robust estimates	1,032	1,037	1,040	1,042	1,044	1,046
Panel (b): In- and out-of-sample distorted JD						
True parameters	951	954	958	960	963	965
(P)MLE	569	709	790	819	853	871
Robust estimates	917	923	929	933	940	945

Notes: The table shows the simulation based average *ex-post* realized utilities in wealth equivalents of an expected utility maximizer ($\vartheta = 0$) and her robust counterpart. Panel (a) depicts realized utilities when the data follow a clean jump-diffusion (JD) process in- and out-of-sample for the three cases of knowing the parameters, estimating them with maximum likelihood, and using robust estimates. Panel (b) shows the same quantities when the data contain small, time-varying distortions to the jump component. Portfolio weights for realized utilities are computed based on five years of daily data generated by the stock process specified in Equation (25). The econometric constant used in the robust estimation is chosen to ensure 95% efficiency in case the true data generating process is a clean jump-diffusion ($c_{Huber} = 40$); $\mu = 0.08$, $\sigma = 0.15$, $\lambda_Y = 3$, $\mu_Y = -0.01$, $\sigma_Y = 0.04$. The number of simulations equals 10,000.

Table 4: Comparative statics equity premium

		BS	JD (with ambiguity aversion)					
ϑ			0	1	2	3	4	5
	0	0.0000	0.0275	0.0316	0.0359	0.0402	0.0446	0.0490
	1	0.0225	0.0503	0.0677	0.0677	0.0677	0.0677	0.0677
γ	2	0.0450	0.0780	0.0826	0.0874	0.0924	0.0975	0.1028
	3	0.0675	0.1108	0.1559	0.2048	0.2579	0.3156	0.3780
	4	0.0900	0.1489	0.3098	0.5059	0.7461	1.0368	1.3786
	5	0.1125	0.1926	0.6201	1.2907	2.3698	3.9401	5.7840

Notes: The table shows the equity premium as a function of risk aversion γ and ambiguity aversion ϑ in a jump-diffusion (JD) economy (Equation ((29))). For comparison, the equity premium implied by a Black-Scholes (BS) world is also given as function of γ (Equation ((28))). To single out the premium dependence on the preference parameters, the stock price process is characterized by $\mu = 0.08$, $\sigma = 0.15$, $\lambda_Y = 3$, $\mu_Y = -0.01$, $\sigma_Y = 0.04$.

Table 5: Accuracy of estimated ambiguity and risk aversion parameters

		JD		Distorted JD	
		(P)MLE	Robust estimates	(P)MLE	Robust estimates
RMSE	γ	7.48	7.74	12.09	9.01
	ϑ	0.41	0.44	0.71	0.53
MAD	γ	3.01	3.15	4.46	3.26
	ϑ	0.17	0.18	0.55	0.30

Notes: The table shows the accuracy of estimated ambiguity and risk aversion parameters based on equilibrium option prices for maximum likelihood and robust parameter estimates. Estimates are obtained from 5 years of daily data following a jump-diffusion (JD) process. Columns 1 and 2 show the root mean squared error (RMSE) and mean absolute deviation (MAD) of estimated ambiguity and risk aversion parameters compared to the true values underlying equilibrium option prices ($\vartheta = 1$, $\gamma = 3$). Columns 3 and 4 show the same quantities when the data contain small, time-varying distortions to the jump component (Equation (25)).

Table 6: Bid-ask spreads due to ambiguity

K	80	90	95	100	105	110	115	120	130	140	150	160	170
Median bid-ask spread													
True parameters	0.8117	1.6380	1.9886	2.2083	2.2664	2.1736	1.9662	1.6910	1.1096	0.6468	0.3491	0.1801	0.0907
Distorted, classical	1.2887	2.4159	2.9083	3.2410	3.3855	3.3486	3.1557	2.8684	2.1554	1.4931	0.9874	0.6380	0.4077
Distorted, robust	0.7757	1.5517	1.8839	2.0902	2.1494	2.0702	1.8859	1.6319	1.0808	0.6385	0.3497	0.1831	0.0933
Clean, classical	0.7509	1.5364	1.8724	2.0870	2.1503	2.0676	1.8745	1.6110	1.0524	0.6191	0.3353	0.1748	0.0882
Clean, robust	0.7548	1.5416	1.8800	2.0946	2.1592	2.0771	1.8832	1.6203	1.0572	0.6221	0.3369	0.1758	0.0887
Median relative bid-ask spread													
True parameters	0.0356	0.1131	0.1806	0.2728	0.3937	0.5485	0.7435	0.9864	1.6667	2.7375	4.4766	7.4046	12.489
Distorted, classical	0.0559	0.1600	0.2456	0.3572	0.4945	0.6590	0.8552	1.0846	1.6588	2.4330	3.4486	4.7933	6.5511
Distorted, robust	0.0340	0.1067	0.1701	0.2569	0.3696	0.5147	0.6981	0.9261	1.5548	2.5047	3.9268	6.1226	9.7053
Clean, classical	0.0330	0.1061	0.1707	0.2591	0.3751	0.5247	0.7118	0.9449	1.6082	2.6719	4.3938	7.2662	12.252
Clean, robust	0.0331	0.1065	0.1713	0.2599	0.3765	0.5260	0.7136	0.9479	1.6110	2.6760	4.3956	7.2646	12.240

Notes: Panel (a) shows median bid-ask spreads for estimates obtained based on a perturbed jump diffusion (Equation (25)). Bid and ask prices are computed from Equations (31) and (32), respectively. Relative bid-ask spreads in Panel (b) are computed as ask minus bid divided by the price of the option without ambiguity. In each panel we consider five parameter cases, namely true parameters (line 1) as well as classical and robust parameter estimates obtained from a clean jump diffusion (lines 4 and 5) as well as a perturbed (Equation (25)) jump diffusion (lines 2 and 3). The time to maturity of the option is 1 year, $S(0) = 100$, and $r = 3\%$. The entropy bound is $\eta = 0.00061$, which corresponds to the entropy of a jump diffusion with a jump standard deviation of $2\sigma_Y = 0.08$. The number of simulations equals 10,000.

Table 7: Parameter estimates

	Panel (a): Discrete time (1929–2009)						
	d_1	d_2	d_3	d_4	$Var(v_1)$	$Var(v_2)$	$Corr(v_1, v_2)$
Classical	0.0137	−0.0365	0.0032	0.9892	0.0030	0.0037	−0.9493
	(0.0106)	(0.0124)	(0.0035)	(0.0040)	(0.0001)	(0.0001)	(0.0214)
Robust	0.0260	−0.0337	0.0063	0.9905	0.0020	0.0026	−0.9549
	(0.0131)	(0.0147)	(0.0040)	(0.0045)	(0.0001)	(0.0001)	(0.0412)
	Panel (b): Continuous time (1929–2009)						
	σ	α	ν	κ_X	$\bar{x}$	b	ρ
Classical	0.1890	0.1800	0.0390	0.1305	-3.3768	0.2130	-0.9495
	(0.0019)	(0.1283)	(0.0414)	(0.0003)	(31.8639)	(0.0024)	(0.0213)
Robust	0.1570	0.3223	0.0765	0.1149	-3.5418	0.1770	-0.9552
	(0.0032)	(0.1587)	(0.0482)	(0.0004)	(34.5276)	(0.0036)	(0.0409)

Notes: Panel (a) shows parameter estimates for the discrete time VAR model in Equation (35). The econometric constant c_{Huber} used in the robust estimation equals 6 to achieve 95% efficiency if the data actually have been generated by the reference model. Asymptotic standard errors are shown in parenthesis. Panel (b) depicts the corresponding continuous time parameter estimates and standard errors (Equations (36) and (38)).

Table 8: Portfolio weights and in-sample utility

Panel (a): Portfolio weights with classical estimates

ϑ		0	1	2	3	4	5
	2	71.77%	48.30%	36.11%	28.83%	23.99%	20.55%
	5	29.53%	24.48%	20.90%	18.24%	16.17%	14.53%
γ	10	14.82%	13.43%	12.28%	11.31%	10.48%	9.76%
	15	9.90%	9.26%	8.69%	8.20%	7.75%	7.35%
	20	7.43%	7.06%	6.73%	6.43%	6.15%	5.90%

Panel (b): Portfolio weights with robust estimates

ϑ		0	1	2	3	4	5
	2	89.78%	80.70%	68.35%	56.42%	47.22%	40.46%
	5	58.98%	49.09%	41.85%	36.43%	32.24%	28.91%
γ	10	30.11%	27.19%	24.78%	22.77%	21.06%	19.59%
	15	20.13%	18.78%	17.60%	16.56%	15.63%	14.81%
	20	15.11%	14.34%	13.64%	13.01%	12.43%	11.90%

Panel (c): Realized wealth with classical estimates

ϑ		0	1	2	3	4	5
	2	108.77	107.23	106.35	105.82	105.47	105.22
	5	105.87	105.50	105.24	105.05	104.90	104.78
γ	10	104.80	104.70	104.62	104.55	104.49	104.43
	15	104.44	104.40	104.36	104.32	104.29	104.26
	20	104.26	104.24	104.21	104.19	104.17	104.15

Panel (d): Realized wealth with robust estimates

ϑ		0	1	2	3	4	5
	2	109.52	109.37	108.82	108.06	107.40	106.90
	5	108.22	107.54	107.00	106.59	106.26	106.00
γ	10	106.10	105.87	105.68	105.52	105.38	105.27
	15	105.31	105.20	105.11	105.03	104.96	104.89
	20	104.91	104.85	104.80	104.75	104.70	104.66

Notes: Portfolio allocation as function of risk aversion (γ) and ambiguity aversion (ϑ). Panels (a) and (b) show the time series mean of the optimal allocation into equities for both classical and robust parameter estimates, respectively. In each month t, π^* is computed with an investment horizon of $T = 1$ year using the parameter estimates from Table 7 and the current dividend-price ratio. Panels (c) and (d) show the average realized wealth from these investment policies.

Table 9: Sensitivity of portfolio weights to return perturbations

	Portfolio weights			Realized wealth		
	Min	Max	Range	Min	Max	Range
Perturbation of 3 most influential returns by up to $\pm 30\%$						
Classical, $\vartheta = 0$	27.68%	31.39%	3.71%	104.58	104.96	0.3879
Robust, $\vartheta = 0$	58.23%	59.58%	1.34%	105.95	106.05	0.0908
Classical, $\vartheta = 3$	17.18%	19.30%	2.13%	104.25	104.49	0.2344
Robust, $\vartheta = 3$	35.91%	36.81%	0.90%	105.20	105.25	0.0548
Perturbation of 3 most influential returns by up to $\pm 10\%$						
Classical, $\vartheta = 0$	28.86%	30.18%	1.32%	104.69	104.83	0.1371
Robust, $\vartheta = 0$	58.23%	59.58%	1.34%	105.97	106.02	0.0442
Classical, $\vartheta = 3$	17.85%	18.61%	0.75%	104.32	104.41	0.0826
Robust, $\vartheta = 3$	35.91%	36.81%	0.90%	105.20	105.25	0.0528
Perturbation of 6 most influential returns by up to $\pm 30\%$						
Classical, $\vartheta = 0$	24.87%	35.48%	10.62%	104.22	105.50	1.2741
Robust, $\vartheta = 0$	57.91%	60.86%	2.94%	105.76	106.31	0.5495
Classical, $\vartheta = 3$	15.52%	21.64%	6.11%	104.03	104.82	0.7885
Robust, $\vartheta = 3$	35.67%	37.75%	2.08%	105.05	105.44	0.3899
Perturbation of 6 most influential returns by up to $\pm 10\%$						
Classical, $\vartheta = 0$	27.71%	31.45%	3.74%	104.54	105.00	0.4593
Robust, $\vartheta = 0$	57.91%	60.14%	2.22%	105.89	106.08	0.1967
Classical, $\vartheta = 3$	17.19%	19.32%	2.12%	104.23	104.51	0.2775
Robust, $\vartheta = 3$	35.67%	37.17%	1.50%	105.13	105.28	0.1505

Notes: The table shows the effect of slightly perturbing the most extreme return observations on the optimal portfolio weights and the corresponding realized wealths. Weight sensitivities are reported for the standard CRRA agent ($\vartheta = 0$) and her robust counterpart ($\vartheta = 3$) for classical and robust parameter estimates. Continuously compounded annual risk-free interest rate $r_f = 3.59\%$, coefficient of risk aversion $\gamma = 5$, initial log dividend-price ratio equals $\log(3.66\%)$. The econometric constant c_{Huber} used in the estimation of CRSP market returns between 1929 and 2009 equals 6 to achieve 95% efficiency if the data actually have been generated by the reference model.

Table 10: Realized out-of-sample utility

		Classical estimates			Robust estimates		
ϑ		0	1	3	0	1	3
				Dividend yield as predictor			
γ	5	106.13	105.96	105.76	107.00	106.67	106.28
	15	105.48	105.46	105.43	105.78	105.74	105.67
				Random walk			
γ	5	105.56	105.49	105.41	107.04	106.73	106.34
	15	105.29	105.28	105.27	105.79	105.75	105.68
				Mixed strategy (5% conf.)			
γ	5	105.56	105.49	105.41	106.84	106.55	106.20
	15	105.29	105.28	105.27	105.72	105.69	105.62
				Mixed strategy (10% conf.)			
γ	5	105.56	105.49	105.41	107.41	107.01	106.53
	15	105.29	105.28	105.27	105.91	105.86	105.78
				Mixed strategy (20% conf.)			
γ	5	105.84	105.72	105.58	107.50	107.09	106.59
	15	105.39	105.37	105.35	105.94	105.89	105.81
				Mixed strategy (30% conf.)			
γ	5	106.09	105.93	105.74	107.48	107.07	106.58
	15	105.47	105.45	105.42	105.94	105.89	105.80

Notes: The table shows the implied out-of-sample utility (wealth equivalent) for an agent using the dividend yield as predictor, the random walk model and a mixed strategy based on evidence for predictability. Utility is reported for different levels of risk aversion (γ) and ambiguity aversion (ϑ) for classical and robust parameter estimates. In each month t, π^* is computed with an investment horizon of $T = 1$ year using parameter estimates based on information up to time t and the current dividend-price ratio. In the mixed strategy cases, the agent tests in each month t for the presence of predictability using the null hypothesis $H_0 : \nu = 0$ against $H_A : \nu > 0$ in Equation (37) at different confidence levels. π^* is then either computed based on a random walk model or based on the optimal policy laid out in Equation (40). The data extends from 1929–2009; the initial sample size for the estimation is 25 years. The econometric constant c_{Huber} used in the estimation equals 6 to achieve 95% efficiency if the data actually have been generated by the reference model.

Table 11: Realized out-of-sample utility in different subsamples

ϑ	0	3	0	3	0	3
	1954–1972		1973–1990		1991–2009	
DP, Classical	105.10	104.58	109.05	108.69	104.24	104.01
DP, Robust	106.12	105.19	110.99	109.85	103.91	103.79
RW, Classical	104.65	104.29	108.07	108.09	103.97	103.85
RW, Robust	107.54	106.10	108.04	108.07	105.57	104.85
Mixed 5%, Classical	104.65	104.29	108.07	108.09	103.97	103.85
Mixed 5%, Robust	107.54	106.10	108.75	108.49	104.26	104.01
Mixed 10%, Classical	104.65	104.29	108.07	108.09	103.97	103.85
Mixed 10%, Robust	107.54	106.10	110.77	109.70	103.93	103.81
Mixed 20%, Classical	104.65	104.29	108.74	108.50	104.15	103.95
Mixed 20%, Robust	107.62	106.14	110.99	109.85	103.90	103.79
Mixed 30%, Classical	104.65	104.29	109.35	108.88	104.29	104.04
Mixed 30%, Robust	107.55	106.10	110.99	109.85	103.91	103.79

Notes: The table shows the implied out-of-sample utility (wealth equivalent) for different subsamples. Utility is reported for the standard CRRA agent ($\vartheta = 0$) and her robust counterpart ($\vartheta = 3$) for classical and robust parameter estimates. The coefficient of risk aversion γ equals 5. In the mixed strategy cases, the agent tests in each month t for the presence of predictability using the null hypothesis $H_0 : \nu = 0$ against $H_A : \nu > 0$ in Equation (37) at different confidence levels. π^* is then either computed based on a random walk model (no predictability) or based on the optimal policy laid out in Equation (40) (in case of predictability), both using an investment horizon of $T = 1$ year and parameter estimates incorporating information up to time t and the current dividend-price ratio. The data extends from 1929–2009; the initial sample size for the estimation is 25 years. The econometric constant c_{Huber} used in the estimation equals 6 to achieve 95% efficiency if the data actually have been generated by the reference model.

References

Amihud, Y., C. M. Hurvich, and Y. Whang, 2009, "Multiple-Predictor Regressions: Hypothesis Testing", *The Review of Financial Studies*, 22(1), 413–434.

Anderson, E., L. Hansen, and T. Sargent, 2000, "Robustness, Detection and the Price of Risk", Unpublished manuscript, Stanford University.

Barberis, N., 2000, "Investing for the Long Run when Returns are Predictable", *The Journal of Finance*, 55(1), 225–264.

Benigno, P., and S. Nisticò, 2009, "International Portfolio Allocation under Model Uncertainty", NBER Working paper 14734, University of Rome Tor Vergata.

Bergstrom, A. R., 1984, "Continuous Time Stochastic Models and Issues of Aggregation Over Time", in *Handbook of Econometrics*, ed. by Z. Griliches, and M. D. Intriligator. Elsevier, Amsterdam, The Netherlands, vol. 2, chap. 20, pp. 1145–1212.

Björk, T., and I. Slinko, 2006, "Towards a General Theory of Good-Deal Bounds", *Review of Finance*, 10(2), 221–260.

Boyarchenko, N., 2010, "Ambiguity Shifts and the 2007 Crisis", Working paper, University of Chicago.

Boyle, P., S. Feng, W. Tian, and T. Wang, 2008, "Robust Stochastic Discount Factors", *The Review of Financial Studies*, 21(3), 1077–1122.

Bustos, O., and V. Yohai, 1986, "Robust Estimates for ARMA Models", *Journal of the American Statistical Association*, 81(393), 155–168.

Caballero, R., and A. Krishnamurthy, 2008, "Collective Risk Management in a Flight to Quality Episode", *The Journal of Finance*, 63(5), 2195–2230.

Campbell, J., G. Chacko, J. Rodriguez, and L. Viceira, 2004, "Strategic Asset Allocation in a Continuous-Time VAR Model", *Journal of Economic Dynamics and Control*, 28(11), 2195–2214.

Campbell, J., and M. Yogo, 2006, "Efficient Tests of Stock Return Predictability", *Journal of Financial Economics*, 81(1), 27–60.

Cao, H., T. Wang, and H. Zhang, 2005, "Model Uncertainty, Limited Market Participation, and Asset Prices", *The Review of Financial Studies*, 18(4), 1219–1251.

Carr, P., D. B. Madan, and J. J. V. Alvarez, 2010, "Markets, Profits, Capital, Leverage and Return", Working paper, Morgan Stanley and University of Maryland.

Chen, H., N. Ju, and J. Miao, 2009, "Dynamic Asset Allocation with Ambiguous Return Predictability", Working paper, Boston University.

Chen, L., 2009, "On the Reversal of Return and Dividend Growth Predictability: A Tale of Two Periods", *Journal of Financial Economics*, 92(1), 128–151.

Chen, Z., and L. Epstein, 2002, "Ambiguity and Asset Returns in Continuous Time", *Econometrica*, 70(4), 1403–1443.

Cochrane, J., and J. Saá-Requejo, 2000, "Beyond arbitrage: Good-Deal Asset Price Bounds in Incomplete Markets", *Journal of Political Economy*, 108(1), 79–119.

Dell'Aquila, R., E. Ronchetti, and F. Trojani, 2003, "Robust GMM Analysis of Models for the Short Rate Process", *Journal of Empirical Finance*, 10(3), 373–397.

Drechsler, I., 2009, "Uncertainty, Time-Varying Fear, and Asset Prices", Working paper, University of Pennsylvania.

Easley, D., and M. O'Hara, 2009, "Ambiguity and Nonparticipation: The Role of Regulation", *The Review of Financial Studies*, 22(5), 1817–1843.

Ellsberg, D., 1961, "Risk, Ambiguity, and the Savage Axioms", *The Quarterly Journal of Economics*, 75(4), 643–669.

Epstein, L., and J. Miao, 2003, "Two-Person Dynamic Equilibrium under Ambiguity", *Journal of Economic Dynamics and Control*, 27(7), 1253–1288.

Epstein, L., and M. Schneider, 2003, "Recursive Multiple Priors", *Journal of Economic Theory*, 113(1), 32–50.

——— , 2007, "Learning Under Ambiguity", *Review of Economic Studies*, 74(4), 1275–1303.

——— , 2008, "Ambiguity, Information Quality, and Asset Pricing", *The Journal of Finance*, 63(1), 197–228.

——— , 2010, "Ambiguity and Asset Markets", *Annual Review of Financial Economics*, 2, 315–346.

Epstein, L., and T. Wang, 1994, "Intertemporal Asset Pricing Under Knightian Uncertainty", *Econometrica*, 62(2), 283–322.

Fama, E., and K. French, 1988, "Dividend Yields and Expected Stock Returns", *Journal of Financial Economics*, 22(1), 3–25.

Gagliardini, P., P. Porchia, and F. Trojani, 2009, "Ambiguity Aversion and the Term Structure of Interest Rates", *The Review of Financial Studies*, 22(10), 4157–4188.

Gagliardini, P., F. Trojani, and G. Urga, 2005, "Robust GMM Tests for Structural Breaks", *Journal of Econometrics*, 129(1–2), 139–182.

Genton, M., and E. Ronchetti, 2003, "Robust Indirect Inference", *Journal of the American Statistical Association*, 98(461), 67–76.

Gilboa, I., and D. Schmeidler, 1989, "Maxmin Expected Utility with Non-unique Prior", *Journal of Mathematical Economics*, 18(2), 141–153.

Hampel, F., 1974, "The Influence Curve and its Role in Robust Estimation", *Journal of the American Statistical Association*, 69(346), 383–393.

Hampel, F., E. Ronchetti, P. Rousseeuw, and W. Stahel, 1986, *Robust Statistics: The Approach Based on Influence Functions.* Wiley, New York, USA.

Hansen, L., and T. Sargent, 2001, "Robust Control and Model Uncertainty", *American Economic Review*, 91(2), 60–66.

Huber, P. J., 1964, "Robust Estimation of a Location Parameter", *The Annals of Mathematical Statistics*, 35(1), 73–101.

——— , 1981, *Robust Statistics.* Wiley, Hoboken, USA.

Ju, N., and J. Miao, 2010, "Ambiguity, Learning, and Asset Returns", Working paper, Boston University.

Kim, T., and E. Omberg, 1996, "Dynamic Nonmyopic Portfolio Behavior", *The Review of Financial Studies*, 9(1), 141–161.

Klibanoff, P., M. Marinacci, and S. Mukerji, 2005, "A Smooth Model of Decision Making Under Ambiguity", *Econometrica*, 73(6), 1849–1892.

——— , 2009, "Recursive Smooth Ambiguity Preferences", *Journal of Economic Theory*, 144(3), 930–976.

Knight, F., 1921, *Risk, Uncertainty and Profit*. Houghton-Mifflin, Boston, USA.

Krishnakumar, J., and E. Ronchetti, 1997, "Robust Estimators for Simultaneous Equations Models", *Journal of Econometrics*, 78(2), 295–314.

Künsch, H., 1984, "Infinitesimal Robustness for Autoregressive Processes", *Annals of Statistics*, 12(3), 843–863.

La Vecchia, D., and F. Trojani, 2010, "Infinitesimal Robustness for Diffusions", *Journal of the American Statistical Association*, 105(490), 703–712.

Leippold, M., and F. Trojani, 2008, "Learning and Asset Prices under Ambiguous Information", *The Review of Financial Studies*, 21(6), 2565–2597.

Lettau, M., and S. C. Ludvigson, 2010, "Measuring and Modeling Variation in the Risk-Return Tradeoff", in *Handbook of Financial Econometrics*, ed. by Y. Ait-Sahalia, and L. P. Hansen. Elsevier, Amsterdam, The Netherlands, vol. 1, chap. 11.

Lettau, M., and S. Van Nieuwerburgh, 2008, "Reconciling the Return Predictability Evidence", *The Review of Financial Studies*, 21(4), 1607–1652.

Liu, J., J. Pan, and T. Wang, 2005, "An Equilibrium Model for Rare-event Premia and Its Implication for Option Smirks", *The Review of Financial Studies*, 18(1), 131–163.

Maenhout, P., 2004, "Robust Portfolio Rules and Asset Pricing", *The Review of Financial Studies*, 17(4), 951–983.

——— , 2006, "Robust Portfolio Rules and Detection-Error Probabilities for a Mean-Reverting Risk Premium", *Journal of Economic Theory*, 128(1), 136–163.

Mancini, L., E. Ronchetti, and F. Trojani, 2005, "Optimal Conditionally Unbiased Bounded-Influence Inference in Dynamic Location and Scale Models", *Journal of the American Statistical Association*, 100(470), 628–641.

Mancini, L., and F. Trojani, 2010, "Robust Value at Risk Prediction", *Journal of Financial Econometrics*, forthcoming.

Merton, R. C., 1969, "Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case", *The Review of Economics and Statistics*, 51(3), 247–257.

——— , 1976, "Option Pricing when Underlying Stock Returns are Discontinuous", *Journal of Economic Dynamics and Control*, 3, 125–144.

Miao, J., and N. Wang, 2010, "Risk, Uncertainty, and Option Exercise", *Journal of Economic Dynamics and Control*, forthcoming.

Muler, N., and V. J. Yohai, 1999, "Robust Estimates for ARCH Processes", *Journal of Time Series Analysis*, 23(3), 341–375.

Ortelli, C., and F. Trojani, 2005, "Robust efficient method of moments", *Journal of Econometrics*, 127(1), 69–97.

Ronchetti, E., and F. Trojani, 2001, "Robust Inference with GMM Estimators", *Journal of Econometrics*, 101(1), 37–69.

Routledge, B., and S. Zin, 2009, "Model Uncertainty and Liquidity", *Review of Economic Dynamics*, 12(4), 543–566.

Sakata, S., and H. White, 1998, "High Breakdown Point Conditional Dispersion Estimation with Application to S & P 500 Daily Returns Volatility", *Econometrica*, 66(3), 529–567.

Trojani, F., and A. Sbuelz, 2008, "Asset Prices with Locally-Constrained-Entropy Recursive Multiple Priors Utility", *Journal of Economic Dynamics and Control*, 32(4), 3695–3717.

Trojani, F., and P. Vanini, 2002, "A Note on Robustness in Merton's Model of Intertemporal Consumption and Portfolio Choice", *Journal of Economic Dynamics and Control*, 28(2), 1079–1113.

——— , 2004, "Robustness and Ambiguity Aversion in General Equilibrium", *Review of Finance*, 8(2), 279–324.

Ulrich, M., 2009, "Inflation Ambiguity and the Term Structure of Arbitrage-Free US Government Bonds", Working paper, Columbia University.

Uppal, R., and T. Wang, 2003, "Model Misspecification and Underdiversification", *The Journal of Finance*, 58(6), 2465–2486.

Welch, I., and A. Goyal, 2008, "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction", *The Review of Financial Studies*, 21(4), 1455–1508.

Xia, Y., 2001, "Learning About Predictability: The Effects of Parameter Uncertainty on Dynamic Asset Allocation", *The Journal of Finance*, 56(1), 205–246.

Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums

Loriano Mancini, Angelo Ranaldo[16]*, and Jan Wrampelmeyer*

This paper was presented at:

- 8th Swiss Doctoral Workshop in Finance, Gerzensee, Switzerland (June 2009)
- CEPR/Study Center Gerzensee European Summer Symposium in Financial Markets, Gerzensee, Switzerland (July 2009)
- Warwick Business School FERC 2009 conference on Individual Decision Making, High Frequency Econometrics and Limit Order Book Dynamics, Coventry, UK (September 2009)
- 59th Midwest Finance Association Annual Meeting, Las Vegas, USA (February 2010)
- 46th Eastern Finance Association Annual Meeting, Miami, USA (April 2010)
- Workshop on International Asset Pricing at the University of Leicester, UK (June 2010)
- 5th Annual Meeting of the Swiss Finance Institute, Zurich, Switzerland (November, 2010)

The most recent version of the paper is available at `http://ssrn.com/abstract=1447869`

Abstract

We utilize ultra-high frequency data to accurately measure liquidity in the foreign exchange (FX) market, quantify the amount of commonality in liquidity across different exchange rates, and provide evidence for liquidity risk being a risk factor for carry trade returns. FX liquidity exhibits significant cross-sectional and temporal variation during the financial crisis of 2007–2009. As sudden shocks to market-wide liquidity have important implications for regulators and investors, liquidity is decomposed into an idiosyncratic and a common component. Empirical results show that liquidity comoves strongly across currency pairs and that systematic FX liquidity decreases dramatically during the financial crisis. Consistent with the theory of liquidity spirals, we document that FX market liquidity is related to funding liquidity and liquidity of equity markets. In an asset pricing context, we introduce a tradable FX liquidity risk factor which is shown to account for most of the cross-sectional variation in carry trade returns.

---

[16]The views expressed herein are those of the authors and not necessarily those of the Swiss National Bank, which does not accept any responsibility for the contents and opinions expressed in this paper.

1. Introduction

Recent events during the financial crisis of 2007–2009 have highlighted the fact that liquidity is a crucial yet elusive concept in all financial markets. The evaporation of liquidity in the funding and foreign exchange markets prompted policy makers and central banks around the world to implement several unconventional policies in an unprecedented coordinated effort to stabilize the financial system and to restore liquidity. According to the Federal Reserves's chairman Ben Bernanke, "weak liquidity risk controls were a common source of the problems many firms have faced [throughout the crisis]" (Bernanke, 2008). Therefore, measuring liquidity and evaluating exposure to liquidity risk is of relevance not only for investors, but also for central bankers, regulators, as well as academics.

While there exists an extensive literature studying the concept of liquidity in equity markets, liquidity in the foreign exchange (FX) market has mostly been neglected, although the FX market is by far the world's largest financial market. The estimated average daily trading volume of four trillion US dollar in 2010 (Bank for International Settlements, 2010) corresponds to more than ten times that of global equity markets (World Federation of Exchanges, 2009). Due to this size, the FX market is commonly regarded as extremely liquid. Nevertheless, events during the financial crisis and the study on currency crashes by Brunnermeier, Nagel, and Pedersen (2009) highlight the importance of liquidity in the FX market. Similarly, Burnside (2009) argues that liquidity frictions potentially play a crucial role in explaining the profitability of carry trades as "liquidity spirals" or "liquidity black holes" (Morris and Shin, 2004) aggravate currency crashes and pose a great risk to carry traders. In the model of Brunnermeier and Pedersen (2009), an increase in perceived downside risk of carry trades lowers funding liquidity, which in turn leads to a decrease in market liquidity and an unwinding of positions.[1] FX markets are extensively used to fund short-term positions, thus illiquidity in FX markets affects funding costs and increases rollover risks, which also has implications for hedging.

This paper fills the gap in the literature by investigating the issue of liquidity in the FX market empirically, allowing to test the theory of liquidity spirals and to analyze the impact of liquidity risk on carry trade returns. To that end, we (i) accurately measure liquidity in the FX market during the crisis of 2007–2009, (ii) quantify the amount of commonality in liquidity across different exchange rates, (iii) relate FX market liquidity to measures of funding liquidity and liquidity of equity markets, and (iv) provide evidence for liquidity risk being a risk factor for carry trade returns. In addition, our analysis contributes to the growing literature on the financial crisis which tries to understand the main stylized facts to determine the causes of the recent market turmoil.

We compute benchmark FX liquidity on a daily basis using a new comprehensive data set of ultra-high frequency return and order flow data. Ranging from January 2007 to December 2009, our sample includes the financial crisis and is thus highly relevant for analyzing liquidity. By applying a variety of liquidity measures covering the dimensions of price impact, return

[1] Further recent papers analyzing crash risk in currency markets include Jurek (2009), Farhi, Fraiberger, Gabaix, Ranciere, and Verdelhan (2009), as well as Plantin and Shin (2010).

reversal, trading cost, and price dispersion we document time series as well as cross-sectional variation in FX liquidity. For instance, bid-ask spreads surged during the financial crisis on average reaching as much as 19 times the pre-crisis level. The decline in liquidity is most severe for AUD/USD, which is frequently used as an investment currency in carry trades. We quantify the potential cost of illiquidity in the FX markets by a realistic carry trade example which shows that FX illiquidity can aggravate losses during market turmoil by as much as 25%. Thus, our analysis helps investors to better assess the risk and potential losses due to FX exposure. Moreover, it allows central banks to effectively monitor liquidity in FX markets supporting policy decisions. Understanding in detail what happened in the FX market during the crisis is a necessary requirement to be able to investigate what caused the crisis and to optimize future monetary policy.

Liquidity of all exchange rates decreased dramatically during the financial crisis indicating commonality in liquidity across FX rates. Such sudden shocks to market-wide liquidity have important implications for regulators as well as investors. Regulators are concerned about the stability of financial markets, whereas investors worry about the risk-return profile of their asset allocation. Decomposing liquidity into an idiosyncratic and a common component allows investors to exploit portfolio theory to reap diversification benefits with respect to liquidity risk. Therefore, we construct a time series of systematic FX liquidity representing the common component in liquidity across different exchange rates. Empirical results show that liquidity comoves strongly across currencies supporting the notion of liquidity being the sum of a common and an exchange rate specific component. Systematic FX liquidity decreases dramatically during the financial crisis, especially after the default of Lehman Brothers in September 2008.

The finding of strong commonality supports the model of Brunnermeier and Pedersen (2009) which predicts comovement in liquidity of different exchange rates during liquidity spirals. To corroborate this evidence, we relate systematic FX liquidity to proxies for uncertainty as well as funding liquidity in financial markets finding that a decrease in these variables leads to lower FX market liquidity. Moreover, market-wide FX liquidity comoves with equity liquidity which is also consistent with the presence of funding liquidity constraints during the financial crisis. Interestingly, AUD/USD is most sensitive to changes in common liquidity, further highlighting the risk inherent in carry trades.

The last part of the paper investigates whether liquidity risk helps to explain daily variation in carry trade returns. Shocks to common FX liquidity are shown to be persistent and we introduce a tradable liquidity risk factor by constructing a portfolio of carry trades. This novel risk factor is correlated to shocks in liquidity as well as the carry trade risk factor of Lustig, Roussanov, and Verdelhan (2010). Compared to the latter, our liquidity risk factor has a clearer and more direct interpretation following from the theory on liquidity spirals which hypothesizes that a drop in market liquidity triggers large exchange rate movements. Apart from stressing the importance of liquidity risk in the determination of FX returns, this finding supports risk-based explanations for deviations from Uncovered Interest Rate Parity (UIP) as classical tests do not include liquidity risk.

The empirical analysis of this paper is related to the substantial strain of literature dealing with liquidity in equity markets. Motivated by the theoretical model of Amihud and Mendelson (1986), various authors have developed measures of liquidity for different time horizons.[2] Similarly, Chordia, Roll, and Subrahmanyam (2000) as well as Hasbrouck and Seppi (2001) derive measures of systematic liquidity and document that liquidity of individual stocks comoves with industry- and market-wide liquidity. To capture different dimensions of liquidity in a single measure, Korajczyk and Sadka (2008) apply principle component analysis to extract a latent systematic liquidity factor both across stocks as well as across liquidity measures. Recently, these measures of common liquidity have been related to equity returns to assess the existence of a return premium for systematic liquidity risk. By augmenting the Fama and French (1993) three-factor model by a liquidity risk factor, Pástor and Stambaugh (2003) find that aggregate liquidity risk is priced in the cross section of stock returns. The studies by Acharya and Pedersen (2005), Sadka (2006), and Korajczyk and Sadka (2008) lend further support to this hypothesis.

Despite its importance, only very few studies exist on liquidity in the FX market, mainly focusing on the explanation of the contemporaneous correlation between order flow and exchange rate returns documented by Evans and Lyons (2002). Using a unique database from a commercial bank, Marsh and O'Rourke (2005) investigate the effect of customer order flows on exchange rate returns. Based on price impact regressions, the authors show that the correlation between order flow and exchange rate movements varies among different groups of customers, suggesting that transitory liquidity effects do not cause the contemporaneous correlation described by Evans and Lyons (2002). On the contrary, Breedon and Vitale (2010) argue that portfolio balancing temporarily leads to liquidity risk premiums and, therefore, affects exchange rates as long as dealers hold undesired inventory. In line with this result, Berger, Chaboud, Chernenko, Howorka, and Wright (2008) document a prominent role of liquidity effects in the relation between order flow and exchange rate movements in their study of Electronic Brokerage System (EBS) data. However, none of these papers systematically measures benchmark liquidity or investigates commonality in liquidity as is done in this paper.

Burnside, Eichenbaum, Kleshchelski, and Rebelo (2011) stress the profitability of carry trades, finding an average annual excess return close to 5% over the period 1976–2007 for a simple carry trade strategy. Recently, Lustig, Roussanov, and Verdelhan (2010) developed a factor model in the spirit of Fama and French (1993) for foreign exchange returns. They posit that a single carry trade risk factor, which is related to the difference in excess returns for exchange rates with large and small interest rate differentials, is able to explain most of the variation in currency excess returns over uncovered interest rate parity. Menkhoff, Sarno, Schmeling, and Schrimpf (2011) adapt this model by stressing the role of volatility risk. The rationale for in-

[2]Among others, Chordia, Roll, and Subrahmanyam (2001) use trading activity and transaction cost measures to derive daily estimates of liquidity from intraday data. In case only daily data are available, Hasbrouck (2009) estimates the effective cost of trades by relying on the spread model of Roll (1984). Alternatively, Amihud (2002) advocates a measure of illiquidity computed as the average ratio of absolute stock return to its trading volume, which can be interpreted as a proxy of price impact. Pástor and Stambaugh (2003) measure stock market liquidity on a monthly basis based on daily return reversal, summarizing the link between returns and lagged order flow. Goyenko, Holden, and Trzcinka (2009) compare various proxies of liquidity against ultra-high frequency benchmarks.

vestigating excess returns is the plethora of papers which document the failure of UIP, rooted in the seminal works of Hansen and Hodrick (1980) as well as Fama (1984). Hodrick and Srivastava (1986) argue that a time-varying risk-premium which is negatively correlated with the expected rate of depreciation is economically plausible and might help to explain the forward bias. This risk-based explanation for the failure of UIP motivates the study of excess currency returns in an asset pricing context. Engel (1992) argues that the forward exchange rate may include a liquidity premium as well as a risk premium. The paper at hand contributes to this strain of literature by highlighting the role of market frictions and liquidity risk to explain variations in carry trade returns.

The remainder of this paper is structured as follows: The data set and measures of liquidity are presented in Section 2. Liquidity in the FX market is investigated empirically in Section 3. Section 4 introduces measures for systematic liquidity and documents commonality in liquidity between different FX rates. Properties of common liquidity such as the relation to funding liquidity and liquidity of equity markets are discussed in Section 5. Evidence for the importance of a liquidity risk factor for the determination of carry trade returns is presented in Section 6. Section 7 concludes.

2. Measuring Foreign Exchange Liquidity

2.1. The Data Set

Next to the fact that the FX market is less transparent than stock and bond markets, because customers cannot trade on a centralized exchange, the main reason why liquidity in FX markets has not been studied previously in more detail is the paucity of available data. However, in recent years two electronic platforms have emerged as the leading trading systems providing an excellent source of currency trade and quote data. These electronic limit order books match buyers and sellers automatically, leading to the spot interdealer reference price. Via the Swiss National Bank it was possible to gain access to a new data set from EBS including historical data on a one second basis of the most important currency pairs between January 2007 and December 2009. With a market share of more than 60%, EBS has become the leading global marketplace for spot interdealer trading in foreign exchange. For the two most important currency pairs, EUR/USD and USD/JPY, the vast majority of spot trading is represented by the EBS data set (Chaboud, Chernenko, and Wright, 2007). EBS best bid and ask prices as well as volume indicators are available and the direction of trades is known, which is crucial for an accurate estimation of liquidity as it avoids using any Lee and Ready (1991) type rule to infer trade directions. All EBS quotes are transactable, thus, they reliably represent the prevalent exchange rate. Moreover, all dealers on the EBS platform are prescreened for credit, thus, counterparty risk is not a concern when analyzing this data set.[3]

In this paper nine currency pairs will be investigated in detail, namely the AUD/USD, EUR/CHF, EUR/GBP, EUR/JPY, EUR/USD, GBP/USD, USD/CAD, USD/CHF, and USD/

[3]See Chaboud, Chernenko, and Wright (2007) for more information and a descriptive study of the EBS database.

JPY exchange rates. For each exchange rate, the irregularly spaced raw data are processed to construct second-by-second price and volume series, each containing 86,400 observations per day. At every second the midpoint of best bid and ask quotes or the transaction price of deals is used to construct one-second log-returns. For the sake of improved interpretability, these exchange rate returns are multiplied by 10^4 to obtain basis points as the unit of measurement. Observations between Friday 10pm to Sunday 10pm GMT[4] are excluded since only minimal trading activity is observed during these non-standard hours. Moreover, we drop US holidays and other days with unusual light trading activity from the data set.

This ultra-high frequency data set allows for a very accurate estimation of liquidity in the FX market. Goyenko, Holden, and Trzcinka (2009) document the added value of high frequency data when measuring liquidity. For a portfolios of stocks, the time series correlation between sophisticated high frequency liquidity benchmarks and lower frequency proxies (e.g. Roll (1984) or Amihud (2002)) can be as low as 0.018. Even the best proxy (Holden, 2009) achieves only a moderate correlation of 0.62 for certain portfolios. For individual assets these correlations are likely to be even smaller. Thus, when analyzing liquidity it is crucial to rely on high-quality data as we do in this paper.

2.2. Liquidity Measures

This section presents various liquidity measures that we utilize to investigate liquidity in the foreign exchange market. Liquidity is a complex concept with different facets, thus, we classify our measures into three categories, namely price impact and return reversal, trading cost as well as price dispersion.

Price Impact and Return Reversal

The first dimensions of liquidity are the price impact of a trade and the subsequent return reversal. Evans and Lyons (2002) document that contemporaneous order flow is important in the determination of FX returns. Conceptually related to Kyle (1985), the price impact of a trade measures how much the exchange rate changes in response to a given level of order flow. The larger the price impact, the more the exchange rate moves following a trade, reflecting lower liquidity. Moreover, if a currency is illiquid, part of the price impact will be temporary as net buying (selling) pressure leads to an excessive appreciation (depreciation) of the currency followed by a reversal to the fundamental value (Campbell, Grossman, and Wang, 1993). The magnitude of this resilience effect determines the return reversal dimension of liquidity, i.e., the more liquid a currency, the smaller is the temporary price change accompanying order flow.

Our dataset allows for an accurate estimation of price impact and return reversal and does not rely on proxies such as, for instance, the ones proposed by Amihud (2002) and Pástor and Stambaugh (2003). Letting $r_{t_i} = p_{t_i} - p_{t_{i-1}}$, v_{b,t_i}, and v_{s,t_i} denote the log exchange rate return, the volume of buyer initiated trades and the volume of seller initiated trades at time t_i during

[4]GMT is used throughout this paper.

day t, respectively, price impact and return reversal can be modeled as:

$$r_{t_i} = \theta_t + \varphi_t(v_{b,t_i} - v_{s,t_i}) + \sum_{k=1}^{K} \gamma_{t,k}(v_{b,t_{i-k}} - v_{s,t_{i-k}}) + \varepsilon_{t_i}. \tag{1}$$

By estimating the parameter vector $\boldsymbol{\theta}_t = [\theta_t \quad \varphi_t \quad \gamma_{t,1} \ldots \gamma_{t,K}]$ on each day, we are able to directly compute the liquidity dimensions of price impact and return reversal on a daily basis. To ensure that the estimates are not effected by potential outliers, we apply robust regression techniques to estimate the model parameters. The estimation is described in detail in Appendix A. It is expected that the price impact of a trade $L^{(pi)} = \varphi_t$ is positive due to the supply and demand effect of net buying pressure as presented by Evans and Lyons (2002). The overall return reversal is measured by $L^{(rr)} = \gamma_t = \sum_{k=1}^{K} \gamma_{t,k}$, which is expected to be negative.

The intraday frequency to estimate Model (1) should be low enough to distinguish return reversal from simple bid-ask bouncing, hence, one-second data needs to be aggregated. Furthermore, a lower frequency or a longer lag length K have the advantage of capturing delayed return reversal. On the other hand, the frequency should be high enough to accurately measure contemporaneous impact and to obtain an adequate number of observations for each day. The results presented in this paper are mainly based on one-minute data and $K = 5$. Results for different frequencies are similar, suggesting that our results are robust to the choice of sampling frequency over which we aggregate the data. These results are available from the authors upon request.

Interpreting price impact and return reversal coefficients as measures of liquidity is in line with the literature on structural market microstructure models. The price impact can be attributed to relevant private information that is disclosed through the trading process. Following Bjønnes, Osler, and Rime (2008), dealers in the FX market are not equally well informed, because large banks or brokers with the most customer business can observe aggregate order flow that is informative about the ongoing price discovery process in the interdealer market. Thus, asymmetric information might lead to illiquidity in the market as, for instance, a potential seller might be afraid that the buyer has private information. Return reversal effects can arise because dealers require compensation for inventory risk and transaction cost.

Moreover, Model (1) is consistent with recent theoretical models of limit order books. Rosu (2009) develops a dynamic model which predicts that more liquid assets should exhibit smaller spreads and lower price impact. In line with Foucault, Kadan, and Kandel (2005) prices recover quickly from overshooting following a market order if the market is resilient (i.e. liquid). By measuring the relation between returns and lagged order flow Model (1) captures delayed price adjustments due to lower liquidity.

Trading Cost

The second group of liquidity measures covers the cost aspect of illiquidity. In line with the implementation shortfall approach of Perold (1988), the cost of executing a trade can be assessed by investigating bid-ask spreads. A market is regarded as liquid if the proportional bid-ask spread, QS, is low:

$$L^{(ba)} = (P^A - P^B)/P^M, \tag{2}$$

where the superscripts A, B and M indicate the ask, bid and mid quote, respectively. The latter is defined as $P^M = (P^A + P^B)/2$.

In practice trades are not always executed exactly at the posted bid or ask quotes.[5] Instead, deals frequently transact at better prices, deeming quoted spread measures inappropriate for an accurate assessment of execution costs. Therefore, effective costs are computed by comparing transaction prices with the quotes prevailing at the time of execution. The effective spread is defined as:

$$L^{(ec)} = \begin{cases} (P - P^M)/P^M, & \text{for buyer-initiated trades} \\ (P^M - P)/P^M, & \text{for seller-initiated trades,} \end{cases} \tag{3}$$

with P denoting the transaction price. Since our dataset includes quotes and trades we do not have to rely on proxies for the effective spread (e.g. Roll, 1984; Holden, 2009; Hasbrouck, 2009), but can rather compute it directly from observed data. Daily estimates of illiquidity are obtained by averaging the effective cost of all trades that occurred on day t.

Price Dispersion

If markets are volatile, market makers require a higher compensation for providing liquidity due to the additional risk incurred. Therefore, if volatility is high, liquidity tends to be lower and, thus, intraday price dispersion, $L^{(pd)}$ can be used as a proxy for illiquidity; see, e.g., Chordia, Roll, and Subrahmanyam (2000). To that end, we estimate daily volatility from ultra-high frequency intraday data. Given the presence of market frictions, utilizing classic realized volatility (RV) is inappropriate (Aït-Sahalia, Mykland, and Zhang, 2005). Zhang, Mykland, and Aït-Sahalia (2005) developed a nonparametric estimator which corrects the bias of RV by relying on two time scales. This two-scale realized volatility (TSRV) estimator consistently recovers volatility even if the data are subject to microstructure noise.

Latent Liquidity

All previously presented liquidity measures capture different aspects of liquidity. A natural approach to extract the common information across these measures is Principal Component Analysis (PCA). Principal components can then be interpreted as latent liquidity factors for an individual exchange rate. For exchange rate j, all five demeaned and standardized liquidity measures are collected in the $5 \times T$ matrix $\widetilde{\mathbf{L}}_j$, where T is the number of days in our sample. The usual eigenvector decomposition of the empirical covariance matrix is $\widetilde{\mathbf{L}}_j \widetilde{\mathbf{L}}_j' \mathbf{U} = \mathbf{UD}$, where $\mathbf{U}$ is the 5×5 eigenvector matrix, and $\mathbf{D}$ the 5×5 diagonal matrix of eigenvalues. The time series evolution of all five latent factors is given by $\mathbf{U}'\widetilde{\mathbf{L}}_j$, with for instance, the first principal

[5] For instance new traders might come in, executing orders at a better price or the spread might widen if the size of an order is particularly large. Moreover, in some electronic markets traders may post hidden limit orders which are not reflected in quoted spreads.

component corresponding to the largest eigenvalue. Such a decomposition is repeated for each exchange rate to capture the most salient features of liquidity by a few factors.

Trading Activity

As more active markets tend to be more liquid, measures of trading activity such as the number of trades, the trading volume, the percentage of zero return periods, or the average trading interval are frequently used as an indirect measure of liquidity. Unfortunately, the relation between liquidity and trading activity is not unambiguous. Jones, Kaul, and Lipson (1994) show that trading activity is positively related to volatility, which in turn implies lower liquidity. Melvin and Taylor (2009) document a strong increase in FX trading activity during the financial crisis, which they attribute to "hot potato trading" rather than an increase in market liquidity. Moreover, traders apply order splitting strategies to avoid a significant price impact of large trades. Consequently, trading activity is not used as a proxy for FX liquidity in this paper.

3. Liquidity in the Foreign Exchange Market

3.1. Liquidity of Individual Exchange Rates During the Financial Crisis

To begin the analysis, we estimate liquidity measures using the large data set described in the previous section for each trading day. Descriptive statistics for exchange rate returns, order flow, and various liquidity measures are shown in Tables 1–3. Average daily returns in Table 1 reveal that AUD and GBP depreciated, while EUR, CHF and particularly JPY appreciated during the sample period. For USD/CHF and USD/JPY, the average order flow is large and positive, nevertheless, USD depreciated against CHF as well as JPY. In line with expectations, EUR/USD and USD/JPY are traded most frequently while trading activity is the smallest for AUD/USD and USD/CAD.

[Table 1 about here.]

Tables 2 and 3 depict summary statistics of daily estimates for the various liquidity measures. Interestingly, the average return reversal, γ_t, i.e., the temporary price change accompanying order flow, is negative and therefore captures illiquidity. The median is larger than the mean indicating negative skewness in daily liquidity. Depending on the currency pair, one-minute returns are on average reduced by 0.013 to 0.172 basis points if there was an order flow of 1–5 million in the previous five minutes. This reduction is economically significant given the fact that average five-minute returns are virtually zero. In line with the results of Evans and Lyons (2002) as well as Berger, Chaboud, Chernenko, Howorka, and Wright (2008), the trade impact coefficient, φ_t, is positive. Effective costs are smaller than half the bid-ask spread implying significant within quote trading. Annualized foreign exchange return volatility ranges from 5.9% to more than 14%.

Comparing the liquidity estimates across different currencies, EUR/USD is the most liquid exchange rate, which is in line with the perception of market participants and the fact that it has by far the largest market share in terms of turnover (Bank for International Settlements, 2010). On the other hand, the least liquid currency pairs are USD/CAD and AUD/USD. Despite the fact that GBP/USD is one of the most important exchange rates, it is estimated to be rather illiquid, which can be explained by the fact that GBP/USD is mostly traded on Reuters rather than the EBS trading platform (Chaboud, Chernenko, and Wright, 2007). The high liquidity of EUR/CHF and USD/CHF during the sample period might be related to "flight-to-quality" effects due to perceived safe haven properties of the Swiss franc (Ranaldo and Söderlind, 2010) during the crisis.

[Tables 2 and 3 about here.]

Figure 1 shows effective cost as defined in Equation (3) for all currencies in our sample over time. Most exchange rates are relatively liquid and stable at the beginning of the sample. In line with Melvin and Taylor (2009), who identify August 16, 2007 to be the beginning of the crisis in FX markets, liquidity suddenly decreased during the major unwinding of carry trades in August 2007. In the following months liquidity rebounded slightly for most currency pairs before it started a downward trend at the end of 2007. Melvin and Taylor (2009) attribute this decline mainly to changes in risk appetite and commodity related selling of investment currencies causing investors to deleverage by unwinding carry trades. The decrease in liquidity continued after the collapse of Bear Stearns in March 2008. A potential reason for the increase in liquidity during the second quarter of 2008 is that investors believed that the crisis might be over soon and began to invest again in FX markets. Moreover, central banks around the world supported the financial system by a variety of traditional as well as unconventional policy tools. However, in September and October 2008, liquidity suddenly and heavily dropped following the default of Lehman Brothers. This decline reflects the unprecedented turmoil and uncertainty in financial markets caused by the bankruptcy. During 2009 FX liquidity slowly but steadily returned. However, there are also cross-sectional differences in how liquidity reacts to crisis events. For instance, liquidity of AUD/USD drops quicker and more pronounced compared to other exchange rates following the default of Lehman Brothers. Interestingly, the ranking of exchange rates according to liquidity is rather stable over time.

[Figure 1 about here.]

While Figure 1 only shows effective cost, all other measures of liquidity share similar patterns. Indeed the PCA reveals that one single factor can explain up to 78.9% of variation in the liquidity measures for EUR/USD. Table 4 shows the loadings of the first three principle components for all currency pairs. In particular the first two principle components have clear interpretations. The first component, which on average explains 70% of the variation in liquidity measures, loads roughly equally on price impact, bid-ask spread, effective cost, and price dispersion. The loading on return reversal is consistently smaller for all exchange rates. In contrast, the second principle component is dominated by return reversal and accounts for an additional 15% of variation. These factor loading are remarkably similar across exchange rates. The remaining principle components do not have clear interpretations.

[Table 4 about here.]

To summarize, these results suggest that (i) the level of liquidity varies drastically across exchange rates, (ii) liquidities comove strongly across exchange rates, and (iii) the liquidity based ranking of exchange rates is rather stable over time. Before analyzing all of these aspects in more detail, the next subsection highlights the economic significance and the relevance of illiquidity in the FX markets by quantifying potential losses due to illiquidity for foreign exchange investors.

3.2. Quantifying the Impact of Illiquidity on a Foreign Exchange Investor

To quantify the economic importance of costs due to illiquidity in the FX market we analyze a simple concrete carry trade example. Pinning down FX illiquidity cost is a challenging task, for instance, because the carry trades strategy is frequently enhanced with a maturity mismatch, i.e., long term lending is financed by short term borrowing. Moreover, investors have the choice between secured fixed income assets such as repos and more risky unsecured assets such as interbank loans. However, these aspects pertain to the fixed income markets and have no impact on the costs due to illiquidity in the FX market. Therefore, we abstract from these additional costs and focus on the direct effect of FX illiquidity on investors' profits. Moreover, we keep exchange rates as well as interest rates constant, assume that the speculator is not levered, and abstract from all additional costs which might impact carry trade returns. An extension of the example including leverage and additional costs will be discussed below.

Consider a US speculator who wants to engage in the AUD-JPY carry trade. She plans to fund this trade by borrowing the equivalent of one million USD at a low interest rate, 1%, in Japan and invest at the higher interest rate, 7%, in Australia. She institutes the trade by buying AUD and selling JPY versus USD to earn the interest rate differential. Suppose liquidity is high in the FX market, namely bid-ask spreads are small and given by 2.63bps for AUD/USD, 0.90bps for USD/JPY (minimum pre-crisis level from Table 3). If the US speculator unwinds the carry trade under these liquid conditions, the cost due to illiquidity is very small and amounts to 0.0313% of the trading volume or 0.515% of the profit from the investment. Assume now liquidity is low and for various reasons, such as the impossibility to roll over short term positions in fixed income markets or the necessity to repatriate foreign capital to hold liquid USD denominated assets, the speculator is forced to unwind the carry trade when markets are illiquid. If the bid-ask spread for AUD/USD is 54.03 bps, as during the peak of the crisis in October 2008, the cost due to illiquidity of unwinding the position is 10.70% of the profit! Hence, the cost of unwinding the trade are more than 20 times larger than under the liquid scenario.

Now, consider the illiquidity cost in a slightly more realistic example. In times of low liquidity and unwinding of carry trades, funding currencies (JPY in the example) usually appreciate whereas investment currencies (AUD in the example) depreciate; see e.g. Brunnermeier, Nagel, and Pedersen (2009). Carry traders refer to these sudden movements of investment exchange rates as "going up the stairs and coming down with the elevator". Additionally, speculators often use leverage, which further magnifies potential losses. Suppose the US speculator has levered her investment 4:1 and that the Australian dollar depreciates by 8% before the carry trader manages to unwind the position. Such a scenario is realistic given the sharp movements in exchange rates during fall 2008. In this scenario the carry trader has to bear a substantial loss. Without illiquidity cost in FX markets, the speculator loses 2.56% of the carry volume which corresponds to a loss of 10.24% of her capital. This loss is increased by 25% under illiquid FX market conditions resulting in a 12.81% decrease of capital.

All in all, this example shows that illiquidity in the FX market can lead to significant costs when being forced to liquidate a carry trade position. Note that illiquidity does not only affect speculators. Every investor or company that owns assets denominated in foreign currencies is subject to FX illiquidity risk. Moreover, Figure 1 suggests that the phenomenon of diminishing liquidity and the economic significance of FX illiquidity cost is not limited to a particular currency pair, but rather affects all exchange rates. This commonality in FX liquidity will be investigated in the next section.

4. Commonality in Foreign Exchange Liquidity

Testing for commonality in FX liquidity is crucial as shocks to market-wide liquidity have important implications for investors as well as regulators. Documenting such commonality is also a first necessary step before studying whether liquidity risk is a risk factor for carry trade returns. Although commonality in liquidity has been extensively documented, for example, in the stock market, a priori it is unclear whether such commonality is present in the FX market given the largely different characteristics of the two markets. From a theoretical point of view, the model of Brunnermeier and Pedersen (2009) implies that market liquidity includes common components across securities, because the theory predicts a decline in market liquidity when investors funding liquidity diminishes. To test for commonality in the FX market, a time series of systematic liquidity is constructed representing the common component in liquidity across different exchange rates.

4.1. Common Liquidity Across Exchange Rates

Two approaches have been proposed to extract market-wide liquidity: averaging and PCA. For completeness we implement both techniques, but most of the analysis will be based on the latter. In the first approach an estimate for market-wide FX liquidity is computed simply as the cross-sectional average of liquidity at individual exchange rate level. Chordia, Roll, and Subrahmanyam (2000) and Pástor and Stambaugh (2003) use this method for determining aggregate liquidity in equity markets. In our setting, given a measure of liquidity, daily systematic liquidity $L^{(\cdot)}_{M,t}$ can be estimated as:

$$L^{(\cdot)}_{M,t} = \frac{1}{N} \sum_{j=1}^{N} L^{(\cdot)}_{j,t}, \tag{4}$$

where N is the number of exchange rates and $L^{(\cdot)}_{j,t}$ the liquidity of exchange rate j on day t. In order for systematic liquidity to be less influenced by extreme values, a common practice is to rely on a trimmed mean. Therefore, we exclude the currency pairs with the highest and lowest value for $L_{j,t}^{(\cdot)}$ in the computation of $L_{M,t}^{(\cdot)}$.

Instead of averaging, Hasbrouck and Seppi (2001) as well as Korajczyk and Sadka (2008) rely on principle component analysis (PCA) to extract market-wide liquidity. For each exchange rate, a given liquidity measure is standardized by the time series mean and standard deviation of the average of the liquidity measure obtained from the cross section of exchange rates. Then, the first three principle components across exchange rates are extracted for each liquidity measure, with the first principal component representing market-wide liquidity. Unreported factor loadings show that the first principal component loads roughly equally on the liquidity of each exchange rate. Thus, for each liquidity measure, systematic liquidity based on PCA can be interpreted as a level factor which behaves similarly to the trimmed mean in Equation (4).

Systematic FX liquidity based on averaging different measures of liquidity is depicted in Panels (a)–(e) of Figure 2. The sign of each measure is adjusted such that the measure represents liquidity rather than illiquidity, i.e., an increase in the measure is associated with higher liquidity. All measures of market-wide liquidity uniformly indicate a steep decline in liquidity after September 2008 when the default of Lehman Brothers as well as the rescue of American International Group (AIG) took place. The stabilization of liquidity at the end of 2008 might be related to governments' and central banks' efforts to support the financial sector using numerous unconventional policy measures. For instance, central banks instituted swap lines to provide liquidity on a massive scale and the US government initiated the Troubled Asset Relief Program (TARP). Common FX liquidity almost recovered to the pre-Lehman level in the course of 2009.

[Figure 2 about here.]

4.2. Testing for Commonality in FX Liquidity

To formally test for commonality, for each exchange rate j, the time series of liquidity measure $L_{j,t}^{(\cdot)}, \{t = 1, \ldots, T\}$ is regressed on the first three principle components described above. Table 5, which shows the cross-sectional average of the adjusted-R^2, reveals ample evidence of strong commonality. The first principle component explains between 70% and 90% of the variation in daily FX liquidity depending on which measure is used. As additional support, the R^2 increases further when two or three principle components are included as explanatory variables. The reversal measure exhibits the lowest level of commonality. The commonality, already strong at daily frequency, increases even more when aggregating liquidity measures at weekly and monthly horizons.

[Table 5 about here.]

The R^2 statistics are significantly larger than those typically found for equity data and reported, e.g., in Chordia, Roll, and Subrahmanyam (2000), Hasbrouck and Seppi (2001), and Korajczyk and Sadka (2008). This would imply that commonality in the FX market is stronger than in equity markets. However, it remains to be seen whether this phenomenon is specific to our sample period, namely the financial crisis 2007–2009, as comovement in financial assets in general and liquidity in particular is reinforced during crisis periods. Note that, the nature of the FX market with triangular connections between exchange rates does not explain the strong commonality. Repeating the principle component regression analysis based on only the six exchange rates which include the US dollar results in R^2 of the same magnitude lending further support to the presence of strong commonality[6].

4.3. Latent Systematic Liquidity Across Measures

Korajczyk and Sadka (2008) take the idea of using PCA to extract common liquidity one step further by combining the information contained in various liquidity measures. Empirical evidence on commonality and visual inspection of Figure 2 shows that alternative liquidity measures yield qualitatively similar results. Indeed, the smallest correlation between different aggregate liquidity measures is between 0.66 for daily and 0.91 for monthly data. This high correlation is consistent with all measures proxying for the same underlying latent liquidity factor. Unobserved systematic liquidity can be extracted by assuming a latent factor model for the vector of standardized liquidity measures, which can again be estimated using PCA:

$$\widetilde{\mathbf{L}}_t = \boldsymbol{\beta} L_{M,t}^{(pca)} + \boldsymbol{\xi}_t, \tag{5}$$

where $\widetilde{\mathbf{L}}_t = \left[\widetilde{\mathbf{L}}_t^{(pi)}, \widetilde{\mathbf{L}}_t^{(rr)}, \widetilde{\mathbf{L}}_t^{(ba)}, \widetilde{\mathbf{L}}_t^{(ec)}, \widetilde{\mathbf{L}}_t^{(pd)}\right]'$ with $\widetilde{\mathbf{L}}_t^{(\cdot)} = \left[\widetilde{L}_{1,t}^{(\cdot)}, \ldots, \widetilde{L}_{N,t}^{(\cdot)}\right]$ denotes the vector which stacks all five liquidity measures for all exchange rates $j = 1, \ldots, N$. $\boldsymbol{\beta}$ is the matrix of factor loadings and $\boldsymbol{\xi}_t$ represents FX rate and liquidity measure specific shocks on day t.

The first principle component explains the majority of variation in liquidity of individual exchange rates, further substantiating the evidence for commonality. Additionally, this allows us to use the first latent factor as proxy for systematic liquidity, $L_{M,t}^{(pca)}$, which combines the information across exchange rates as well as across liquidity measures. Similar to the simple measures, the sign of the factor is chosen such that it represents liquidity.

Panel (f) of Figure 2 depicts latent systematic liquidity estimated according to Equation (5). The graph resembles the ones obtained by averaging liquidity of individual exchange rates. Again we find that market-wide FX liquidity decreased after the beginning of the subprime crisis, there is a steep decline in the aftermath of the collapse of Lehman Brothers, and liquidity gradually recovered to normal levels in 2009.

5. Properties of FX Liquidity

5.1. Relation to Proxies of Investors' Fear and Funding Liquidity

What are the reasons for the strong decline in FX liquidity during the crisis? This subsection tries to answer this question by empirically investigating the link between funding liquidity in

[6]Detailed results are omitted for brevity, but are available from the authors upon request.

FX market liquidity. The typical starting point of liquidity spirals is an increase of uncertainty in the economy, which leads to a retraction of funding liquidity. Difficulty in securing funding for business activities in turn lowers market liquidity, especially if investors are forced to liquidate positions. This induces prices to move away from fundamentals leading to increasing losses on existing positions and a further reduction of funding liquidity which reinforces the downward spiral.

Figure 3 illustrates latent market-wide FX liquidity over time together with the Chicago Board Options Exchange Volatility Index (VIX) as well as the TED spread. Primarily an index for the implied volatility of S&P 500 options, the VIX is frequently used as a proxy for investors' fear and uncertainty in financial markets. The TED spread is a proxy for the level of credit risk and funding liquidity in the interbank market[7]. During most of the sample, the severe financial crisis is reflected in a TED spread which is significantly larger than its long-run average of 30–50 basis points. In 2009, the TED spread narrows due to the low interest rate environment.

[Figure 3 about here.]

Interestingly, the VIX as well as the TED spread are strongly negatively correlated with FX liquidity (approximately -0.87 and -0.35 for daily latent liquidity) indicating that investors' fear measured by implied volatility of equity options and credit risk has spillover effects to other asset classes. Even when excluding the period after the default of Lehman Brothers, the negative correlations prevail (approximately -0.66 and -0.36 for daily latent liquidity). These findings are in line with the theory of liquidity spirals. In particular after the default of Lehman Brothers the VIX and the TED spread surged while market liquidity declined.

To substantiate the evidence for liquidity spirals we regress latent FX liquidity on the lagged VIX and TED spread; Table 6 shows the results. Both the VIX as well as the TED spread are negatively related to common FX liquidity. Thus, an increase in investors' uncertainty and a reduction of funding liquidity at time $t-1$ are followed by significantly lower FX market liquidity at time t. These effects explain a large part of the variation in systematic FX liquidity (adjusted-R^2 equal to 0.76) and are statistically significant. Changing the specification of the regression model, e.g., by controlling for lagged FX market liquidity does not alter the conclusions.

[Table 6 about here.]

Additional to a reduction in funding liquidity followed by a liquidity spiral, the increasing integration of international financial markets might strengthen the link between the VIX and FX liquidity as, for instance, the default of Lehman Brothers led to severe repercussions in all financial markets. To investigate this issue further, the relation between systematic FX liquidity and liquidity of equity markets is investigated in the next subsection.

[7] An alternative proxy for funding liquidity is the LIBOR-OIS spread. The results based on this proxy are similar and are available from the authors upon request.

5.2. Relation to Liquidity of the US Equity Market

There exists a number of reasons to expect a connection between equity and FX illiquidity: If liquidity dries up in the FX market, which is the world's largest financial market, this is a good indication for a liquidity crisis with effects in all financial markets. Moreover, an interdependence between illiquidity in the two markets is consistent with the interaction of market and funding liquidity during liquidity spirals as described in the previous subsection. Finally, central bank interventions directly impact the FX market, but have severe implications for other markets and the worldwide economy. For instance, if monetary shocks are reflected in FX market illiquidity first, this might influence future equity illiquidity.

To investigate the relation of liquidity in both markets and potential commonality across these asset class, the measures of market-wide FX liquidity presented in the previous section are compared to systematic liquidity of the US equity market. The latter is estimated based on (i) return reversal (Pástor and Stambaugh, 2003) and (ii) Amihud's (2002) measure utilizing return and volume data of all stocks listed at the New York Stock Exchange (NYSE) and the American Stock Exchange (AMEX)[8]. Figure 4 shows a comparison of liquidity in FX and equity markets based on a sample of 36 (24 for equity return reversal) non-overlapping monthly observations.

[Figure 4 about here.]

The results support the notion that liquidity shocks are systematic across asset classes. The correlation between equity and latent FX liquidity is 0.81 and 0.36 depending on which measure is used to obtain equity market liquidity. Similarly, a Spearman's rho of 0.67 and 0.39 indicates comovement, further corroborating the finding of integrated financial markets. The significantly lower correlation between average FX and equity return reversal can be explained by noise inherent in the latter. Compared to Pástor and Stambaugh's (2003) reversal measure for equity markets, aggregate FX return reversal for monthly data is negative over the whole sample. This desirable result might be caused by the fact that the EBS data set includes more accurate order flow data and that Model (1) is estimated robustly at a higher frequency.

5.3. Idiosyncratic Liquidity and Exposure to Systematic Liquidity

Having documented the strong commonality of FX liquidity during the financial crisis, the question arises how the liquidity of individual exchange rates relates to systematic FX liquidity. To analyze the sensitivity of the liquidity of exchange rate j to a change in market-wide liquidity, we regress individual liquidity, $L^{(\cdot)}_{j,t}$, on common liquidity $L^{(\cdot)}_{M,t}$:

$$L^{(\cdot)}_{j,t} = a_j + b_j L^{(\cdot)}_{M,t} + L^{(\cdot)}_{I,j,t}. \tag{6}$$

The sensitivity is captured by the slope coefficient denoted by b_j. For the sake of interpretability, we rely on effective cost as measure of liquidity and exclude exchange rate j in the computation

[8] We thank Ľuboš Pástor for providing current estimates for the equity return reversal factor on his website: `http://faculty.chicagobooth.edu/lubos.pastor/research/liq_data_1962_2008.txt`.

of $L_{M,t}^{(ec)}$. The estimation results in Table 7 show that the liquidity of every exchange rate positively depends on systematic liquidity. Hence, given the evidence on liquidity spirals, all exchange rates are affected by funding liquidity constraints and the resulting downward move in market liquidity. This effect is most pronounced for AUD/USD which exhibits the largest slope coefficient for market-wide liquidity: A one basis point decrease in systematic FX liquidity leads to a 3.14bps drop in the liquidity of AUD/USD. This result is intuitive given the fact that AUD is a frequently used investment currency of carry traders, who experienced severe funding constraints during the recent crisis. On the other hand, the commonly perceived most liquid exchange rate, EUR/USD exhibits the lowest sensitivity.

[Table 7 about here.]

The exposure to common liquidity explains between 75% and 90% of the variation in liquidity of individual exchange rates. The remaining 10% to 25% are due to movements in exchange rate specific liquidity which is represented by the residuals $L_{j,t}^{I}$ in Equation (6). Figure 5 plots this idiosyncratic liquidity over time for each exchange rate j. In August 2007, when the crisis spread to the FX market, a significant unwinding of the JPY-AUD carry trade took place. On that day, the USD dollar gained more than 2.8% against the Australian dollar and depreciated against the Yen by more than 1.5%. This unwinding is not only mirrored in aggregate market liquidity, but also led to a downward spike in idiosyncratic liquidities of these two exchange rates. AUD/USD exhibits a unique reaction to crisis events with liquidity dropping earlier and more significant compared to other FX rates after the default of Lehman Brothers. Also the time series of idiosyncratic liquidities exhibit interesting patterns. In general, exchange rate specific liquidity is much more volatile after the default of Lehman Brothers. Standard deviations computed for different subsamples confirm this observation. Panel (b) of Table 7 shows that volatility almost tripled for many exchange rates after September 2008.

[Figure 5 about here.]

There exist a number of potential explanations for the cross-sectional differences and time series variation in idiosyncratic liquidity components. Events and announcements during the financial crisis had a large impact on financial markets[9]. Non-synchronized central bank interventions and diverging scales and timing of measures to restore market stability might have led to unique patterns in individual exchange rate liquidity. For instance, an institution of swap lines between the FED and the Swiss National Bank to provide USD liquidity is likely to first and foremost impact liquidity of the USD/CHF exchange rate. Similarly, central banks' reserve management during the crisis might have played a role.

Additionally, there exist trading related aspects that impact the liquidity of individual exchange rates. The crisis led to exuberant uncertainty regarding the size and location of losses, hence, due to increasing counterparty risk, more transactions were settled in the spot market rather than the forward market. Moreover, algorithmic traders are particularly active in certain

[9] A full analysis of crisis events and central bank interventions at individual exchange rate level is beyond the scope of this paper, but represents an interesting avenue for further research.

cross rates (Chaboud, Chiquoine, Hjalmarsson, and Vega, 2009) to exploit arbitrage opportunities. The presence of these liquidity providers potentially causes unique exchange rate specific movements in liquidity.

6. Evidence for Liquidity Risk Premiums

6.1. Shocks to Systematic Liquidity

Given the evidence for liquidity spirals leading to strong declines in systematic FX liquidity, the question arises whether investors demand a premium for being exposed to this liquidity risk. A necessary condition for such premiums to exist is that shocks to market-wide liquidity are persistent, i.e., shocks need to have long-lasting effects to significantly impact investors. Figure 6 depicts the autocorrelation functions for different estimates of systematic liquidity. Clearly, all aggregate liquidity proxies exhibit strong autocorrelation. Therefore, a drop in aggregate liquidity is not likely to be reversed quickly and investors who would like to unwind a position cannot rely on markets being liquid in any short time period.

[Figure 6 about here.]

6.2. Carry Trade Returns

To investigate the role of liquidity in cross-sectional asset pricing, daily dollar log-returns are constructed from spot rates in units of foreign currency per USD. Hence, in contrast to the previous analysis, all returns use the USD as the base currency, which allows for better interpretation of the factors. Additional to FX data, interest rates are necessary to construct risk factors and to analyze liquidity risk premiums as well as excess returns over UIP. Thus, similar to Liu and Maynard (2005) the interest rate differential for the various currencies is computed from LIBOR interest rates, which are obtained from Datastream. LIBOR rates are converted to continuously compounded rates to allow for comparison with FX log-returns, which are computed at the same point in time.

Combining these data sets, the variable of interest is the excess return over UIP:

$$r^e_{j,t+1} = i^f_t - i^d_t - \Delta p_{j,t+1}, \tag{7}$$

where i^f_t and i^d_t represent the foreign and domestic interest rates at day t, respectively. $r^e_{j,t+1}$ denotes the excess return of currency pair j at day t from the perspective of US investors. Alternatively, it can be interpreted as the return from a carry trade in which a US investor who borrows at the domestic and invests in the foreign interest rate is exposed to exchange rate risk. For the purpose of the asset pricing study, gross excess returns are used, because excess returns net of bid-ask spreads overestimate the true cost of trading (Gilmore and Hayashi, 2008). Descriptive statistics for exchange rate returns, interest rate differentials as well as excess returns are depicted in Table 8.

[Table 8 about here.]

Panel (a) shows that the annualized returns of individual exchange rates between January 2007 and December 2009 are larger in absolute value compared to the longer sample of Lustig, Roussanov, and Verdelhan (2010). While prior to the default of Lehman Brothers (Panel (b)) the difference in magnitude is rather small, larger average and extremely volatile returns occur after the collapse (Panel (c)). In general, the interest rate differentials are lower in absolute value in the last subsample mirroring the joint efforts of central banks to alleviate the economic downturn by lowering interest rates.

Typical carry trade funding currencies of low interest rate countries (JPY, CHF) have a positive excess return over the whole sample with the appreciation being strongest after September 2008. These appreciation of funding currencies is consistent with deteriorating liquidity and flight-to-quality episodes (Ranaldo and Söderlind, 2010). Immediately after the default of Lehman Brothers investment currencies which are associated with high interest rates (AUD, NZD) depreciated strongly mirroring liquidity spirals and unwinding of carry trades. However, in the course of 2009, these currencies appreciated against the USD overall resulting in an negative excess return of the US dollar. A common explanation for this appreciation of the investment currencies are the relatively worse prospects for the US economy at that time. Moreover, investors might have started to setup carry trades again because the quantitative easing and historically low interest rates in the United States fueled the search for yields and allowed the dollar to be used as funding currency. Moreover, commodity prices increased again in 2009 which supported commodity related currencies such as the Australian dollar.

The crisis led to significant volatility in exchange rates; the standard deviation of daily FX returns doubled for many exchange rates when comparing the samples before and after the default of Lehman. This strong variation and significant excess returns over UIP in combination with the large literature on risk-based explanations of this failure warrants further analysis.

6.3. Excess Returns and Liquidity

Recently, a number of papers has documented common variation in carry trade returns (see, e.g., Lustig, Roussanov, and Verdelhan (2010) and Menkhoff, Sarno, Schmeling, and Schrimpf (2011)). The results from the previous sections suggest that liquidity risk might contribute to this common variation. Indeed, there is a strong relation between carry trade returns and liquidity. Figure 7 depicts the cumulative return of one dollar invested in the AUD/USD carry trade together with liquidity of that exchange rate. The cumulative AUD/USD carry trade return mirrors movements in liquidity. The unwinding of carry trades on August 16, 2007 resulted in a drop in liquidity and a large negative carry trade return. In parallel to diminishing liquidity, carry trade returns were negative in the period after the default of Lehman Brothers, before recovering in the course of 2009.

[Figure 7 about here.]

To corroborate this visual evidence for a connection between liquidity and carry trade returns, Table 9 shows correlations between carry trade returns and FX liquidity, liquidity shocks, and unexpected liquidity shocks. The liquidity level is measured by latent systematic liquidity from Section 4. In line with Pástor and Stambaugh (2003) and Acharya and Pedersen (2005), liquidity shocks and unexpected liquidity shocks are defined as the residuals from an AR(1) model and an AR(2) model fitted to latent systematic liquidity, respectively.

Investment currencies such as AUD and NZD depreciate contemporaneously with a decrease in liquidity. On the other hand, the Japanese yen appreciates which is consistent with Table 8 and the findings regarding liquidity spirals above. The correlation is largest in absolute value for shocks at the monthly horizon. Also at the daily frequency, correlations between liquidity shocks and carry trade returns are between 50% and 100% larger than the connection between the returns and the common liquidity level.

[Table 9 about here.]

This evidence for correlation between excess returns over UIP and (unexpected) changes in liquidity is consistent with liquidity risk being a risk factor for carry trade returns.

6.4. Liquidity Risk Factor

Following the arbitrage pricing theory of Ross (1976), variation in the cross section of returns is assumed to be caused by different exposure to a small number of risk factors. Lustig, Roussanov, and Verdelhan (2010) propose a factor model for excess FX returns including a "dollar risk factor" capturing the average excess return for a US investor and a "carry trade risk factor", which is long the exchange rates with the largest interest rate differential and short the exchange rates with the smallest interest rate differential. The authors find that the latter explains common variation in carry trade return and suggest that this risk factor captures global risk for which carry traders earn a risk premium. A potential drawback of this model is that the notion of global risk is rather abstract and does not allow for a clear economic interpretation. An alternative, more concrete, explanation for the findings of Lustig, Roussanov, and Verdelhan (2010) is carry traders' exposure to liquidity spirals in conjunction with currency crashes. Evidence from the previous sections shows that liquidity is an important determinant of carry trade returns. Therefore, in this section we construct a liquidity risk factor similar to carry trade risk, HML, of Lustig, Roussanov, and Verdelhan (2010).

To that end, we build a portfolio which is long the two most illiquid and short the two most liquid FX rates on each day t and label this liquidity risk factor IML (illiquid minus liquid). Such a tradable risk factor has the advantage that investors can decide to hedge the associated liquidity risk more easily compared to a factor which is constructed from more complicated liquidity risk measures. Panels (a) and (b) of Figure 8 compare IML to a non-tradable risk factor computed as shocks to market-wide latent liquidity. Both graphs exhibit similar patterns and the factors exhibit much larger variation after the default of Lehman Brothers. Moreover, IML is strongly correlated with HML during our sample period (cf. Panels (a) and (c) of

Figure 8). Thus, our liquidity risk factor captures global risk due to liquidity spirals in periods of strong unwinding of carry trade positions. Note that this interpretation is also more direct compared to the factor model including FX volatility risk of Menkhoff, Sarno, Schmeling, and Schrimpf (2011). While increased volatility might be a consequence of liquidity spirals, first order effects are mirrored FX liquidity.

The second risk factor in our model for carry trade returns is the dollar risk factor from Lustig, Roussanov, and Verdelhan (2010), which is constructed as:

$$AER_t = \frac{1}{N}\sum_{j=1}^{N} r^e_{j,t}, \tag{8}$$

and describes the average excess return, i.e., the return for a US investor who goes long in all N exchange rates available in the sample. As shown in Panel (d) of Figure 8 this level risk factor does not exhibit significant variation compared to both HML as well as IML.

[Figure 8 about here.]

Having described foreign exchange risk factors for explaining carry trade returns, we can estimate a factor model to assess the relative importance and cross-sectional differences in exposure to these factors. To that end, we estimate the following asset pricing model on a daily basis:

$$r^e_{j,t} = \alpha_j + \beta_{AER,j} AER_t + \beta_{IML,j} IML_t + \varepsilon_{j,t}, \tag{9}$$

where $\beta_{AER,j}$ and $\beta_{IML,j}$ denote the exposure of the carry trade return j to the market risk factor and liquidity risk factor, respectively. Any abnormal return that is not explained by our FX risk factors is captured by the constant α_j. The regression results are shown in Table 10. No currency pair exhibits a significant α indicating that the model appropriately captures the characteristics of carry trade returns. This is confirmed by the high adjusted-R^2s which range from approximately 0.50 to 0.90. Thus, the majority of the level and variation of carry trade returns can be explained by exposure to our risk factors. In line with Lustig, Roussanov, and Verdelhan (2010), all exchange rates load rather equally on the market risk factor, which therefore helps to explain the average level of carry trade returns. In contrast, IML betas vary substantially across exchange rates. Interestingly, the Swiss franc and the Japanese yen exhibit the largest negative liquidity betas. Thus, an increase in liquidity risk leads to lower returns of the US dollar against these funding currencies so JPY and CHF tend to appreciate when liquidity risk increases. Similarly, investment currencies like AUD and NZD exhibit the greatest positive liquidity beta implying that these currencies depreciate when liquidity risk increases. These results are again in line with the theory of liquidity spirals and match well to the results in Tables 8 and 9.

[Table 10 about here.]

All in all, our findings provide strong evidence for the presence of a liquidity risk factor in FX returns as an alternative explanation for the global risk premium found by Lustig, Roussanov, and Verdelhan (2010). [10] However, in contrast to the notion of global risk our FX liquidity risk factor has a clearer interpretation as it is directly related to the theory of liquidity spirals and currency crashes. Investors are exposed to these spirals and will thus demand a risk premium as compensation for bearing liquidity risk.

7. Conclusion

The recent financial crisis confirms that illiquidity is ubiquitous in all financial markets during distressed market conditions. Contrary to the common perception of the FX market being extremely liquid at all times, this paper shows that liquidity is an important issue in the FX market. Utilizing an ultra-high frequency data set, we estimate various dimensions of liquidity and document interesting cross-sectional and temporal variation in FX liquidity. The importance of liquidity in the FX market is highlighted by the fact that FX liquidity declined severely during the financial crisis of 2007–2009. This result does not only hold for individual exchange rates, but also for market-wide liquidity. Indeed, liquidity in exchange rates can be decomposed into an idiosyncratic and a common component. This high degree of commonality and the correlation between systematic market liquidity and proxies for funding liquidity supports the theory of liquidity spirals in FX markets. The presence of commonality also has important implications for asset pricing as investors are averse to shocks to market-wide liquidity. We show that shocks to FX liquidity are persistent and negatively correlated with carry trade returns. Given this evidence for the presence of liquidity risk premiums, we introduce a novel tradable liquidity risk factor and document that liquidity risk helps to explain the variation in carry trade returns.

These results have several important implications. Monitoring FX liquidity on a daily basis allows central banks and regulatory authorities to evaluate the effectiveness of their policies. FX liquidity as a real time measure of market stress helps to review how federal authorities coped with crisis events and supports future timing decisions for traditional as well as unconventional policy measures. Moreover, understanding the role of liquidity and liquidity risk helps investors to more adequately assess the risk of their international portfolio positions. In light of the potentially enormous losses from currency crashes coinciding with liquidity spirals, this is particularly crucial for investors applying carry trade strategies.

[10] It remains to be investigated with a longer sample whether our liquidity risk factor is also priced in the cross section of carry trade returns. Given the results of Lustig, Roussanov, and Verdelhan (2010) and the significant correlation between IML and HML, this is likely to be the case.

Appendix A. Estimation of Model (1)

The classic choice to estimate Model (1) is ordinary least squares (OLS) regression. However, high frequency data are likely to contain outliers. Unfortunately, classic OLS estimates are adversely affected by these atypical observations which are separated from the majority of the data. In line with this reasoning, Pástor and Stambaugh (2003) warn that their reversal measure can be very noisy for individual securities.

Removing outliers from the sample is not a meaningful solution since subjective outlier deletion or algorithms as described by Brownlees and Gallo (2006) have the drawback of risking to delete legitimate observations which diminishes the value of the statistical analysis. The approach adopted in this paper is to rely on robust regression techniques. The aim of robust statistics is to obtain parameter estimates, which are not adversely affected by the presence of potential outliers (Hampel, Ronchetti, Rousseeuw, and Stahel, 2005).

Expressing Model (1) in shorthand notation, $r_{t_i} = \boldsymbol{\theta}_t \mathbf{x}_{t_i} + \varepsilon_{t_i}$, where $\mathbf{x}_{t_i}$ includes the intercept and contemporaneous as well as lagged order flows, and ε_{t_i} is an error term, robust parameter estimates for day t are the solutions to:

$$\min_{\boldsymbol{\theta}_t} \sum_{i=1}^{I} \rho\left(\frac{\varepsilon_{t_i}(\boldsymbol{\theta}_t)}{\sigma_t}\right), \tag{10}$$

with I denoting the number of intraday observations. Furthermore, in this equation σ_t represents the scale of the error term and $\rho(\cdot)$ is a bisquare function:

$$\rho(y) = \begin{cases} 1 - \left[1 - (y/k)^2\right]^3 & \text{if} \quad |y| \leq k \\ 1 & \text{if} \quad |y| > k \end{cases}. \tag{11}$$

The first order condition for the optimization problem in Equation (10) is:

$$\sum_{i=1}^{I} \rho'\left(\frac{\varepsilon_{t_i}\left(\widehat{\boldsymbol{\theta}}_t\right)}{\widehat{\sigma_t}}\right) \mathbf{x_{t_i}} = \mathbf{0}, \tag{12}$$

where

$$\rho'(y) = \begin{cases} 6y/k^2 \left[1 - (y/k)^2\right]^2 & \text{if} \quad |y| \leq k \\ 0 & \text{if} \quad |y| > k \end{cases}. \tag{13}$$

In the bisquare function the constant $k = 4.685$ ensures 95% efficiency of $\widehat{\boldsymbol{\theta}}_t$ when ε_{t_i} is normally distributed. Computationally, the parameters are found using iteratively reweighed least squares with a weighting function corresponding to the bisquare function in Equation (11) and an initial estimate for the residual scale of $\widehat{\sigma} = \frac{1}{0.675} \text{median}_{i=1}^{I} \left(|\varepsilon_{t_i}| \, | \, \varepsilon_{t_i} \neq 0\right)$.

Compared to standard OLS, by construction, robust regression estimates are less influenced by potential contamination in the data. Furthermore, standard errors of the robust estimates are typically smaller as outliers inflate classic OLS confidence intervals (Maronna, Martin, and Yohai, 2006).

Appendix B. Figures and Tables

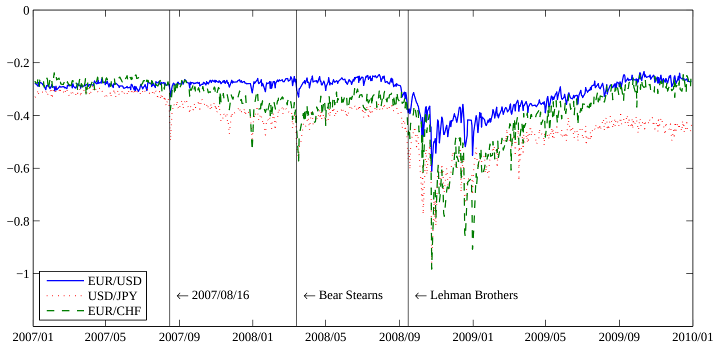

(a)

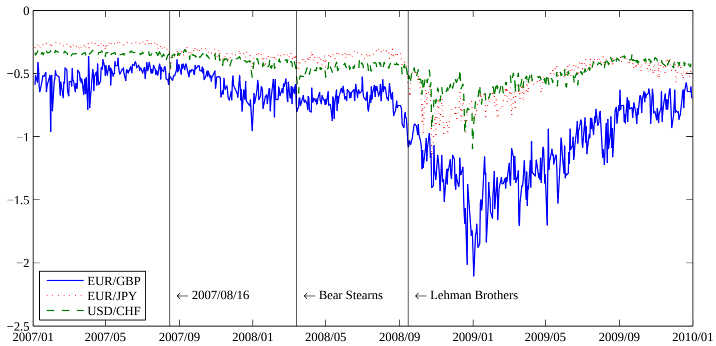

(b)

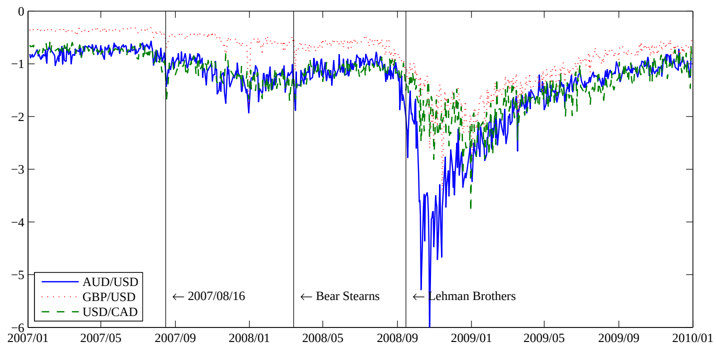

(c)

Figure 1: Daily liquidity estimates based on effective cost. Panel (a) depicts effective cost over time for the estimated most liquid exchange rates (EUR/USD, USD/JPY, and EUR/CHF); Panel (b) shows effective cost for intermediate liquidity currency pairs (EUR/GBP, EUR/JPY, and USD/CHF), whereas the time series of effective cost for the most illiquid currencies (AUD/USD, GBP/USD, and USD/CAD) are plotted in Panel (c). The effective cost is computed as $(P - P^M)/P^M$ for buyer-initiated trades, or $(P^M - P)/P^M$ for seller-initiated trades, where P denotes the transaction price and P^M the mid quote price. The sample is January 2, 2007 – December 30, 2009.

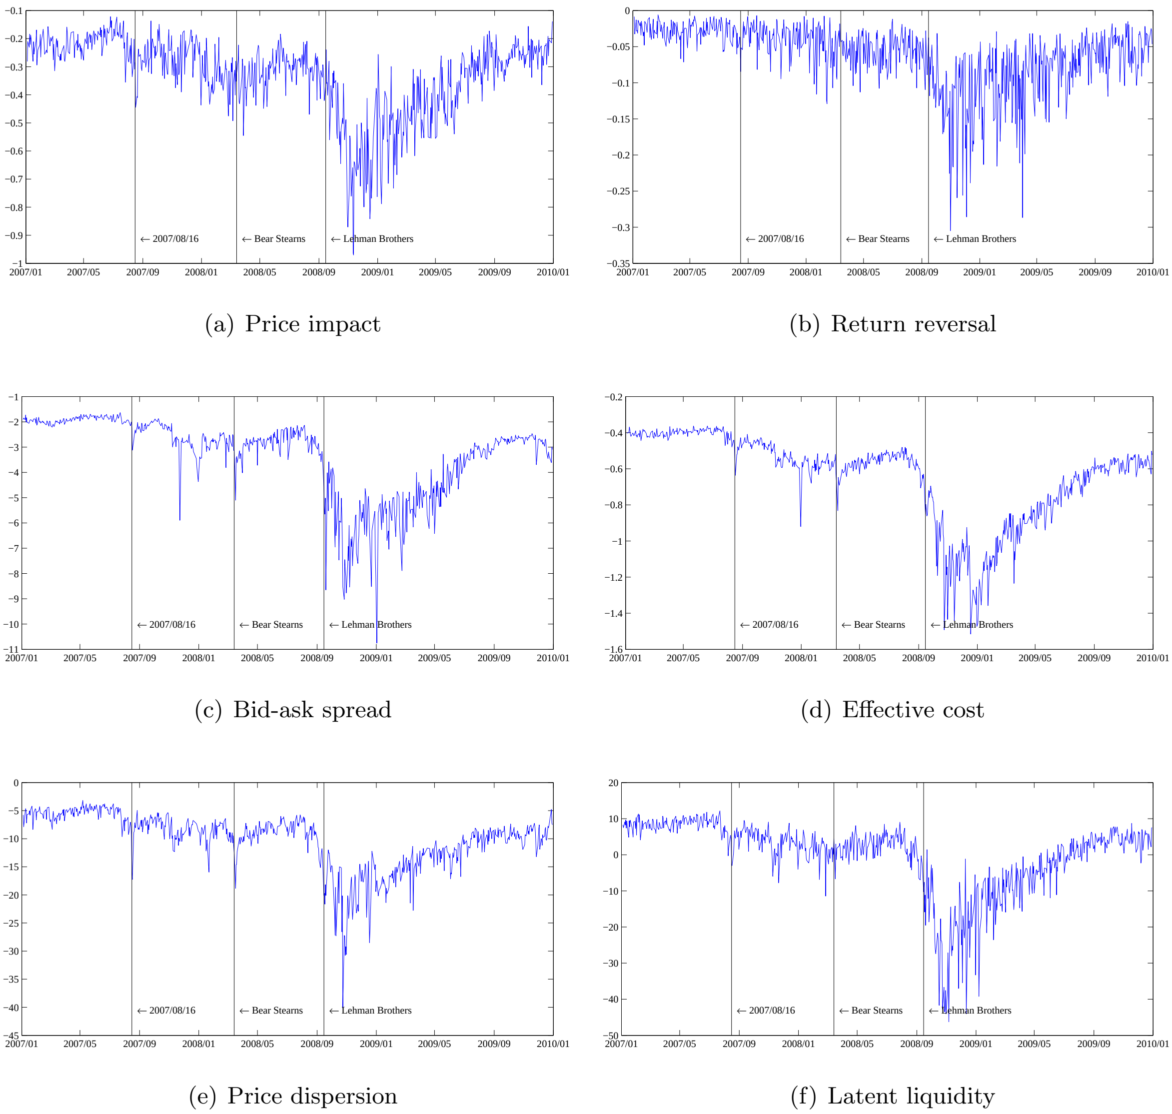

(a) Price impact

(b) Return reversal

(c) Bid-ask spread

(d) Effective cost

(e) Price dispersion

(f) Latent liquidity

Figure 2: Systematic FX liquidity. Panels (a)–(e) depict market-wide FX liquidity based on (within measures) averaging of individual exchange rate liquidity (Equation (4)). Latent systematic liquidity obtained from principle component analysis across exchange rates as well as across liquidity measures (Equation (5)) is depicted in Panel (f). The sign of each liquidity measure is adjusted such that the measure represents liquidity rather than illiquidity. The sample is January 2, 2007 – December 30, 2009.

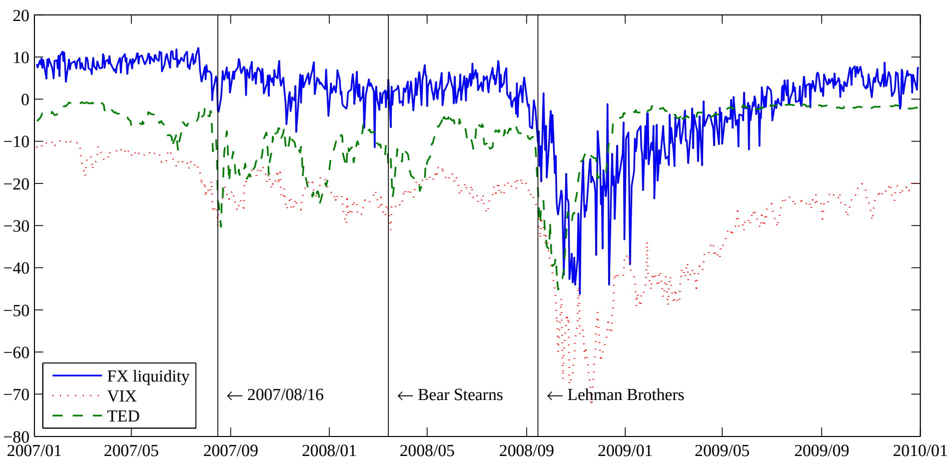

Figure 3: Latent systematic FX liquidity (Equation (5)), the negative of the Chicago Board Options Exchange Volatility Index (VIX) as well as the TED spread. The sample is January 2, 2007 – December 30, 2009.

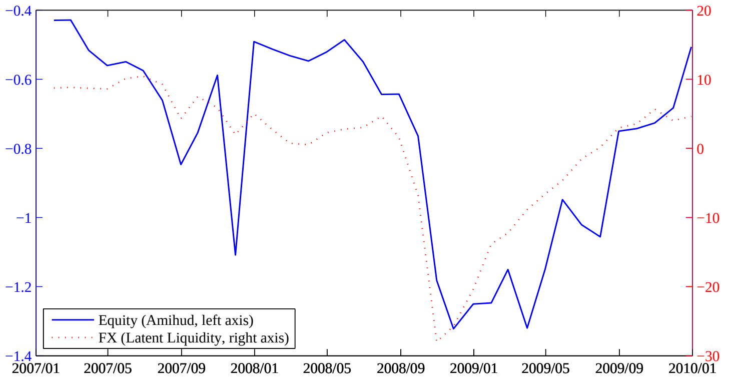

(a)

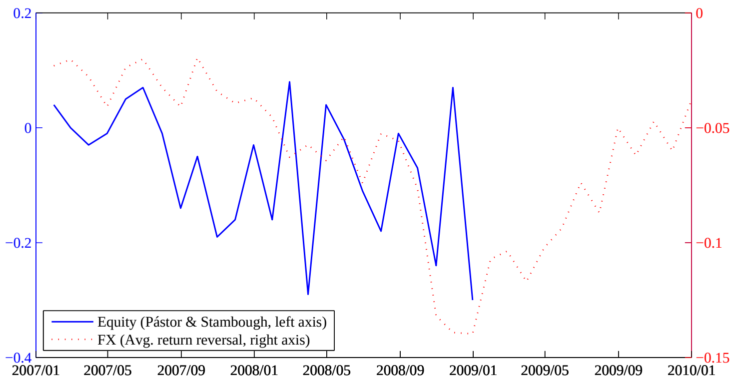

(b)

Figure 4: Non-overlapping monthly systematic FX liquidity and US equity liquidity (estimated from stocks listed on the NYSE and AMEX). In Panel (a), latent FX liquidity obtained from PCA across different liquidity measures is plotted together with Amihud's (2002) measure of equity liquidity. Panel (b) shows the average FX return reversal obtained from Model (1) and equity return reversal (Pástor and Stambaugh, 2003). Each observation t represents estimated liquidity for a given month. Daily FX liquidity is averaged to obtain monthly estimates. The sample is January 2007 – December 2009.

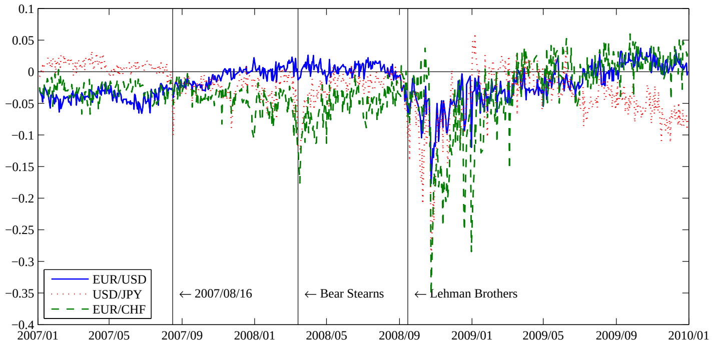

(a)

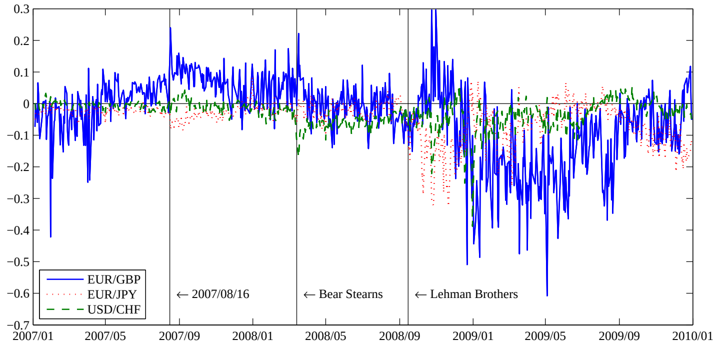

(b)

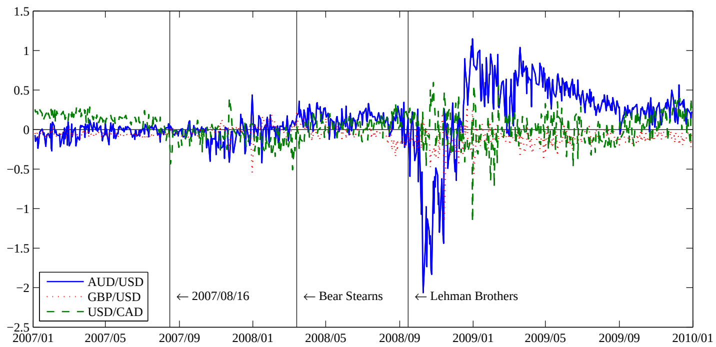

(c)

Figure 5: Idiosyncratic liquidity. Idiosyncratic liquidity of FX rate j is estimated as the residuals from regressing liquidity of exchange rate j on average FX market-wide liquidity (Equation (6)). Both individual exchange rate as well as market-wide liquidity are estimated based on effective cost. The sample is January 2, 2007 – December 30, 2009.

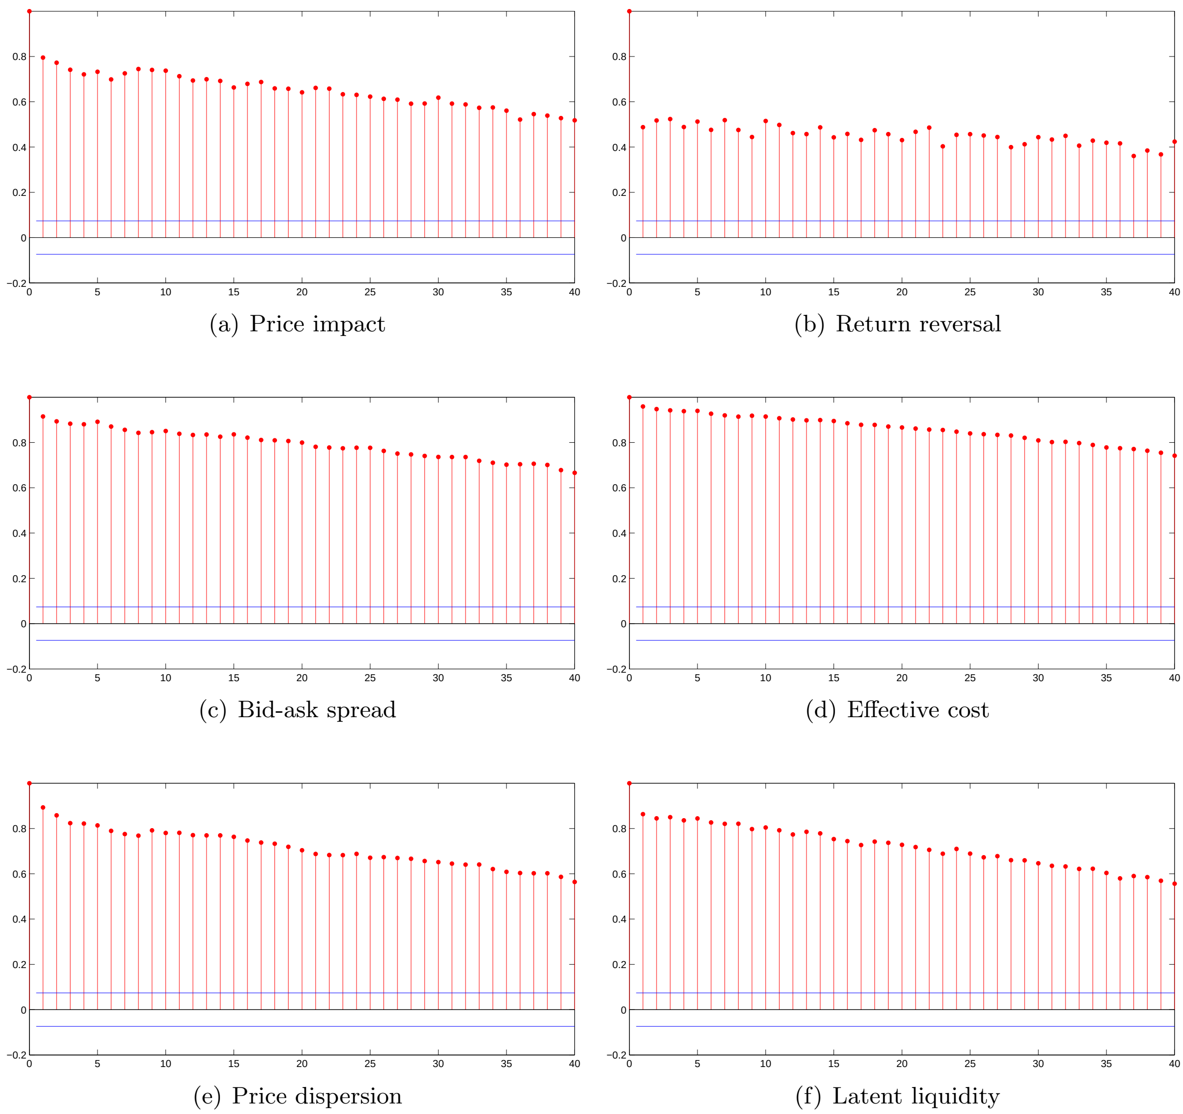

Figure 6: Autocorrelations of daily systematic liquidity. Panels (a)–(e) depict autocorrelations (up to 40 lags) for daily systematic FX liquidity based on (within measures) averaging of individual exchange rate liquidity (Equation (4)). The autocorrelations for latent systematic liquidity obtain from principle component analysis across exchange rates as well as across liquidity measures (Equation (5)) are depicted in Panel (f). The solid horizontal lines indicate upper and lower 95% confidence bounds. The sample is January 2, 2007 – December 30, 2009.

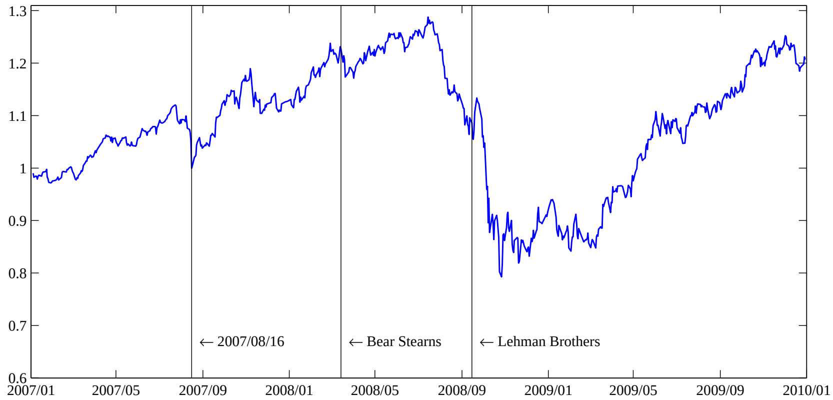

(a) Cumulative AUD/USD carry trade returns

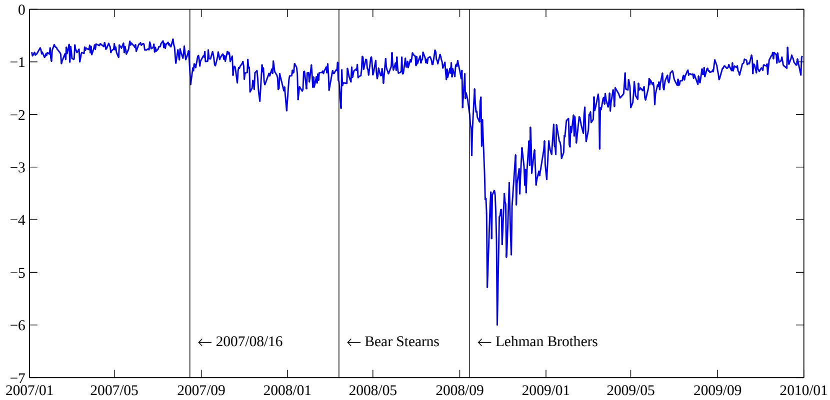

(b) AUD/USD liquidity

Figure 7: Carry trade returns and liquidity. Panel (a) depicts the cumulative return of investing one dollar in the AUD/USD carry trade. AUD/USD liquidity measured by effective cost is shown in Panel (b). The sample is January 2, 2007 – December 30, 2009.

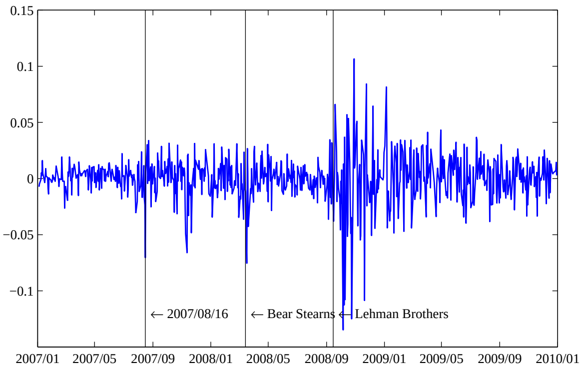

(a) Tradable liquidity factor (IML)

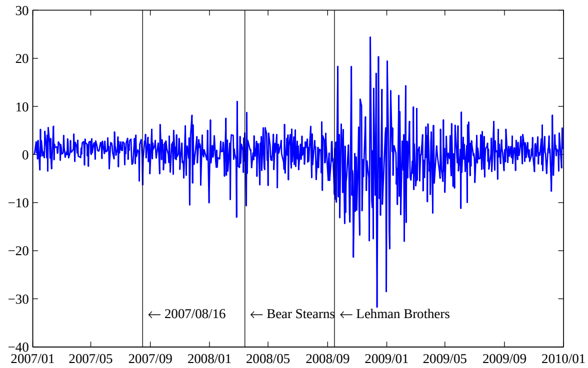

(b) Shocks to latent market-wide FX liquidity

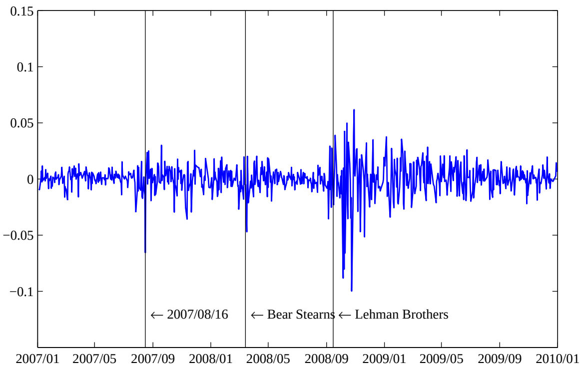

(c) "Carry trade risk factor" (HML)

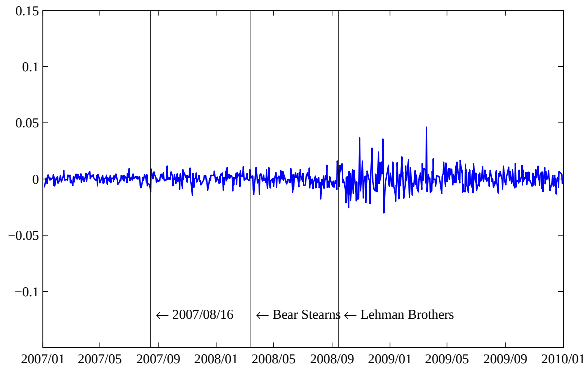

(d) Average excess return (AER)

Figure 8: Time series of risk factors for carry trade returns. The tradable liquidity risk factor, IML, is shown in Panel (a). This factor is defined as the excess return of a portfolio which is long the two most illiquid and short the two most liquid exchange rates. Panel (b) shows shocks to latent systematic FX liquidity (residuals from AR(1) model fitted to $L_{M,t}^{(pca)}$ from Equation (5)). The slope or carry trade risk factor, HML, of Lustig, Roussanov, and Verdelhan (2010), defined as the excess return of a portfolio which is long the two exchange rates with largest interest rate differential and short the two exchange rates with the smallest interest rate differential, is shown in Panel (c). Panel (d) depicts the market risk factor, AER, which is constructed as the average excess return from investing in an equally weighted portfolio of foreign currencies from the perspective of a US investor. The sample is January 2, 2007 – December 30, 2009.

Table 1: Daily FX data

	AUD/USD	EUR/CHF	EUR/GBP	EUR/JPY	EUR/USD	GBP/USD	USD/CAD	USD/CHF	USD/JPY
Whole Sample									
Return (in bps)	0.05	−0.51	3.07	−1.88	0.30	−2.21	−1.04	−1.68	−1.47
Order flow	−9.96	−29.66	−3.89	−31.47	−156.69	5.62	−12.57	59.77	−3.92
# of trades	500	3,751	479	6,750	19,306	794	323	5,605	13,867
Minimum FX rate	0.6021	1.4417	0.6553	114.87	1.2454	1.3773	0.9204	0.9849	86.41
Maximum FX rate	0.9795	1.6797	0.9760	169.51	1.5998	2.1076	1.3015	1.2537	123.92
Prior to default of Lehman Brothers									
Return (in bps)	0.59	0.38	2.30	0.40	1.48	−1.18	−2.39	−1.55	−1.09
Order flow	−14.00	−33.07	−4.86	−50.40	−61.06	1.80	−15.05	38.79	30.56
# of trades	411	4,096	430	7,417	17,575	868	334	6,111	15,013
Minimum FX rate	0.7703	1.5492	0.6553	150.24	1.2893	1.7502	0.9204	0.9849	97.32
Maximum FX rate	0.9795	1.6797	0.8162	169.51	1.5998	2.1076	1.1846	1.2537	123.92
After default of Lehman Brothers									
Return (in bps)	−0.67	−1.67	4.10	−4.89	−1.25	−3.56	0.74	−1.85	−1.97
Order flow	−4.62	−25.17	−2.61	−6.49	−282.89	10.66	−9.29	87.44	−49.42
# of trades	617	3,295	544	5,871	21,590	697	310	4,938	12,355
Minimum FX rate	0.6021	1.4417	0.7737	114.87	1.2454	1.3773	1.0248	0.9963	86.41
Maximum FX rate	0.9369	1.5988	0.9760	156.15	1.5134	1.8560	1.3015	1.2260	107.41

Notes: This table shows mean return (measured in basis points), order flow (measured in volume indicators; a size indicator of 1 corresponds to an order flow of 1–5 million, 6 corresponds to an order flow of 6–10 million, etc.) and number of trades per day for various exchange rates. Moreover, the minimum and maximum exchange rates during the sample are shown. The sample is January 2, 2007 – December 30, 2009.

Table 2: Daily liquidity measures from Model (1)

	AUD/USD	EUR/CHF	EUR/GBP	EUR/JPY	EUR/USD	GBP/USD	USD/CAD	USD/CHF	USD/JPY
				Price impact					
Mean	1.0646	0.1183	0.4970	0.2558	0.0701	0.4322	0.8392	0.1771	0.1123
Median	0.8700	0.0975	0.4045	0.2188	0.0601	0.3344	0.7660	0.1586	0.1030
Std. dev.	0.7748	0.0739	0.2937	0.1462	0.0371	0.3114	0.4692	0.0781	0.0524
% positive	100.00%	99.32%	100.00%	100.00%	100.00%	100.00%	98.91%	100.00%	100.00%
% pos & significant	99.86%	99.05%	99.59%	100.00%	100.00%	99.59%	97.95%	100.00%	100.00%
				Return reversal ($K = 1$)					
Mean	−0.0592	−0.0100	−0.0443	−0.0216	−0.0059	−0.0510	−0.0601	−0.0082	−0.0101
Median	−0.0324	−0.0071	−0.0335	−0.0160	−0.0049	−0.0348	−0.0526	−0.0069	−0.0090
Std. dev.	0.1713	0.0160	0.0804	0.0276	0.0058	0.0778	0.1440	0.0191	0.0090
% negative	67.12%	77.63%	73.12%	82.95%	92.36%	80.49%	68.76%	69.17%	91.68%
% neg & significant	26.33%	40.11%	33.15%	43.52%	63.03%	46.79%	33.56%	27.69%	64.26%
				Return reversal ($K = 3$)					
Mean	−0.1346	−0.0195	−0.0766	−0.0413	−0.0107	−0.0855	−0.1173	−0.0191	−0.0192
Median	−0.0974	−0.0141	−0.0622	−0.0318	−0.0087	−0.0639	−0.0970	−0.0158	−0.0171
Std. dev.	0.2706	0.0247	0.1167	0.0436	0.0093	0.1150	0.2159	0.0320	0.0144
% negative	73.53%	85.68%	77.76%	88.68%	94.95%	84.17%	72.58%	76.40%	96.73%
% neg & significant	48.43%	64.67%	51.98%	69.30%	84.31%	65.48%	52.93%	52.25%	87.59%
				Return reversal ($K = 5$)					
Mean	−0.1717	−0.0255	−0.0929	−0.0531	−0.0132	−0.0980	−0.1596	−0.0255	−0.0243
Median	−0.1244	−0.0186	−0.0786	−0.0410	−0.0106	−0.0748	−0.1308	−0.0218	−0.0205
Std. dev.	0.3240	0.0299	0.1394	0.0516	0.0109	0.1345	0.2605	0.0386	0.0177
% negative	77.22%	88.40%	77.90%	92.36%	96.04%	83.63%	77.76%	78.17%	98.64%
% neg & significant	54.71%	73.26%	57.16%	78.85%	87.86%	68.49%	58.12%	59.62%	91.68%

Notes: This table shows summary statistics for various daily measures of liquidity. Price impact is the robustly estimated coefficient of contemporaneous order flow, φ_t, in a regression of one-minute returns on contemporaneous and lagged order flow (Equation (1)). Return reversal is the sum of the coefficients of lagged order flow, $\sum_{k=1}^{K} \gamma_{t,k}$, in the same regression. The sample is January 2, 2007 – December 30, 2009.

Table 3: Daily liquidity measures

	AUD/USD	EUR/CHF	EUR/GBP	EUR/JPY	EUR/USD	GBP/USD	USD/CAD	USD/CHF	USD/JPY
Bid-ask spread (in bps)									
Mean	5.75	2.07	4.75	2.21	1.05	6.16	8.27	2.50	1.50
Median	4.44	1.81	3.85	1.94	0.91	3.40	6.62	2.28	1.39
Std. dev.	3.87	1.03	2.96	0.96	0.29	7.44	7.63	1.11	0.41
Min	2.64	0.97	1.92	0.98	0.78	1.43	2.88	1.22	0.90
Max	54.03	8.13	29.38	11.49	2.52	67.32	135.72	16.07	3.34
Effective cost (in bps)									
Mean	1.38	0.36	0.81	0.43	0.31	0.81	1.26	0.45	0.42
Median	1.15	0.33	0.70	0.37	0.28	0.65	1.17	0.42	0.41
Std. dev.	0.78	0.11	0.33	0.17	0.06	0.48	0.46	0.11	0.10
Min	0.57	0.24	0.36	0.23	0.23	0.30	0.59	0.30	0.29
Max	6.00	0.98	2.10	1.17	0.61	3.40	3.81	1.10	0.96
Volume weighted effective cost (in bps)									
Mean	1.11	0.28	0.71	0.33	0.21	0.66	1.07	0.34	0.27
Median	0.93	0.25	0.61	0.29	0.19	0.52	0.99	0.31	0.26
Std. dev.	0.64	0.10	0.30	0.14	0.04	0.41	0.41	0.10	0.07
Min	0.28	0.14	0.31	0.15	0.16	0.17	0.39	0.21	0.16
Max	4.93	0.78	1.89	0.89	0.42	2.81	3.33	0.90	0.65
Price dispersion (TSRV, five minutes, in %, annualized)									
Mean	14.25	5.36	8.28	12.26	8.91	11.31	11.84	9.81	10.41
Median	11.67	4.41	7.06	10.22	7.66	8.57	10.99	8.90	9.54
Std. dev.	9.59	3.21	4.36	7.39	4.42	8.29	5.38	4.14	4.84
Min	4.11	1.25	2.56	3.03	2.66	2.90	4.01	3.31	3.13
Max	90.21	29.28	31.52	65.14	29.39	69.05	56.24	33.73	53.36

Notes: This table shows summary statistics for various daily measures of liquidity. Bid-ask spread denotes the average proportional bid-ask spread computed using intraday data for each trading day. Effective cost is the average difference between the transaction price and the bid/ask quote prevailing at the time of the trade. Price dispersion for each trading day is estimated using two-scale realized volatility (TSRV). It is expressed in percentage on an annual basis. The sample is January 2, 2007 – December 30, 2009.

Table 4: Principle component loadings for individual exchange rates

	AUD/USD	EUR/CHF	EUR/GBP	EUR/JPY	EUR/USD	GBP/USD	USD/CAD	USD/CHF	USD/JPY	Average
First principle component loadings										
Return reversal	0.2773	0.3244	0.2443	0.3429	0.3034	0.3662	0.2522	0.1794	0.3534	0.2937
Price impact	0.4304	0.4712	0.4287	0.4687	0.4823	0.4495	0.3876	0.4762	0.4726	0.4519
Bid-ask spread	0.4790	0.4802	0.4531	0.4616	0.4839	0.4224	0.4380	0.4918	0.4808	0.4657
Effective cost	0.5116	0.4883	0.5406	0.4897	0.4703	0.5079	0.5488	0.5195	0.4929	0.5077
Price dispersion	0.4967	0.4513	0.5090	0.4581	0.4691	0.4770	0.5414	0.4788	0.4214	0.4781
Cum. % explained	70.05%	77.21%	59.14%	77.31%	78.90%	63.91%	51.59%	69.18%	77.06%	69.37%
Second principle component loadings										
Return reversal	0.9323	0.9335	0.9476	0.9304	0.9505	0.7661	0.7457	0.9778	0.9269	0.9012
Price impact	−0.3324	−0.0367	0.0677	−0.1601	−0.1175	0.2946	0.4251	−0.0033	−0.0719	0.0073
Bid-ask spread	−0.0812	−0.1643	−0.2312	−0.0919	−0.1539	−0.4061	−0.4311	−0.1287	−0.1905	−0.2088
Effective cost	−0.1100	−0.1876	−0.1511	−0.1513	−0.2066	−0.1134	−0.0256	−0.0786	−0.2054	−0.1366
Price dispersion	−0.0410	−0.2549	−0.1454	−0.2783	−0.1281	−0.3853	−0.2770	−0.1457	−0.2391	−0.2105
Cum. % explained	86.39%	90.72%	76.95%	89.82%	92.98%	78.78%	71.12%	87.70%	89.02%	84.83%
Third principle component loadings										
Return reversal	−0.2171	−0.1288	0.2020	−0.1057	0.0488	−0.4521	−0.6138	0.0973	−0.0991	−0.1410
Price impact	−0.7726	0.4352	−0.8699	−0.4118	−0.1384	0.5145	0.6558	−0.7549	0.3994	−0.1047
Bid-ask spread	0.5254	0.3030	0.3966	0.7258	0.3399	−0.6420	−0.4004	−0.0219	0.3615	0.1764
Effective cost	0.0012	0.1324	0.0807	0.2382	0.5251	0.3370	0.1769	0.0868	0.0495	0.1809
Price dispersion	0.2827	−0.8274	0.1970	−0.4855	−0.7663	0.0720	−0.0390	0.6424	−0.8352	−0.1955
Cum. % explained	94.72%	95.52%	88.39%	95.72%	96.40%	89.17%	86.61%	93.78%	96.96%	93.03%

Notes: This table shows principle component loadings for the first three factors together with the cumulative variation in liquidity that is explained by each factor. For each exchange rate j, all five demeaned and standardized liquidity measures (price impact, return reversal, bid-ask spread, effective cost, price dispersion) are collected in the $5 \times T$ matrix $\widetilde{\mathbf{L}}_j$, where T is the number of sample days. The eigenvector decomposition of the empirical covariance matrix is $\widetilde{\mathbf{L}}_j \, \widetilde{\mathbf{L}}_j' \, \mathbf{U} = \mathbf{UD}$, where $\mathbf{U}$ is the 5×5 eigenvector matrix, and $\mathbf{D}$ the 5×5 diagonal matrix of eigenvalues in descending order. The first three principal component loadings are given by the first three columns of $\mathbf{U}$. The sample is January 2, 2007 – December 30, 2009.

Table 5: Commonality in liquidity using within measure PCA factors

Measure	*Factor 1*	*Factors 1,2*	*Factors 1,2,3*
	Daily data		
Price impact	0.6275	0.7251	0.8080
Return reversal ($K = 1$)	0.2700	0.4011	0.5132
Return reversal ($K = 3$)	0.2961	0.4198	0.5352
Return reversal ($K = 5$)	0.2908	0.4117	0.5234
Bid-ask spread	0.6889	0.7761	0.8553
Effective cost	0.8797	0.9181	0.9405
Effective cost, volume-weighted	0.8877	0.9251	0.9460
Price dispersion (TSRV, one minute)	0.7951	0.8500	0.8991
Price dispersion (TSRV, five minute)	0.8033	0.8563	0.9022
	Weekly data		
Price impact	0.7343	0.8261	0.8877
Return reversal ($K = 1$)	0.4513	0.5982	0.6916
Return reversal ($K = 3$)	0.4972	0.6406	0.7217
Return reversal ($K = 5$)	0.5159	0.6400	0.7269
Bid-ask spread	0.7851	0.8613	0.9090
Effective cost	0.9066	0.9408	0.9608
Effective cost, volume-weighted	0.9197	0.9531	0.9693
Price dispersion (TSRV, one minute)	0.8707	0.9127	0.9402
Price dispersion (TSRV, five minute)	0.8712	0.9191	0.9490
	Monthly data		
Price impact	0.8063	0.8935	0.9441
Return reversal ($K = 1$)	0.6682	0.7933	0.8483
Return reversal ($K = 3$)	0.6962	0.7981	0.8611
Return reversal ($K = 5$)	0.7246	0.8077	0.8630
Bid-ask spread	0.8633	0.9349	0.9648
Effective cost	0.9217	0.9542	0.9740
Effective cost, volume-weighted	0.9361	0.9664	0.9818
Price dispersion (TSRV, one minute)	0.9120	0.9436	0.9641
Price dispersion (TSRV, five minute)	0.9121	0.9471	0.9720

Notes: For each daily standardized measure of liquidity the first three common factors are extracted using principle component analysis. Then, for each exchange rate and each standardized liquidity measure, liquidity is regressed on its common factors. The table shows the average adjusted-R^2 of these regressions using one, two and three factors. The sample is January 2, 2007 – December 30, 2009.

Table 6: Evidence for liquidity spirals in the FX market

	$const$	$L^{pca}_{M,t-1}$	VIX_{t-1}	TED_{t-1}	VXY_{t-1}	Adj. R^2
Coefficient	18.9412		−0.691	−1.2629		0.7645
Std. error	(0.9882)		(0.0403)	(0.4463)		
Coefficient	18.6197		−0.7186			0.7557
Std. error	(1.0376)		(0.0465)			
Coefficient	3.8506			−4.7109		0.1408
Std. error	(1.0768)			(1.3463)		
Coefficient	10.968	0.4182	−0.3977	−0.8077		0.8035
Std. error	(1.3795)	(0.0754)	(0.0523)	(0.2776)		
Coefficient	10.3535	0.4399	−0.3996			0.8001
Std. error	(1.3842)	(0.0773)	(0.0534)			
Coefficient	0.7044	0.8388		−0.8626		0.7503
Std. error	(0.2070)	(0.0243)		(0.2673)		
Coefficient	14.857	0.3569	−0.207	−1.5442	−0.7001	0.8154
Std. error	(1.5976)	(0.0763)	(0.0573)	(0.3371)	(0.1383)	
Coefficient	22.6857		−0.3693	−2.1847	−0.9636	0.7889
Std. error	(1.2624)		(0.0645)	(0.4935)	(0.1897)	
Coefficient	14.3187	0.4324		−1.9598	−1.0817	0.8085
Std. error	(1.6649)	(0.0664)		(0.3594)	(0.1306)	
Coefficient	25.0788			−3.3639	−1.8993	0.7617
Std. error	(1.3440)			(0.5872)	(0.1020)	

Notes: Regression of daily latent systematic FX liquidity ($L^{(pca)}_{M,t}$) on lagged VIX and TED spread. Ten different specifications of the regression model are estimated. The last four specifications additionally control for the JP Morgan Implied Volatility Index for the G7 currencies, VXY. Heteroscedasticity and autocorrelation (HAC) robust standard errors are shown in parenthesis. The sample is January 2, 2007 – December 30, 2009.

Table 7: Sensitivity to changes in common liquidity

	AUD/USD	EUR/CHF	EUR/GBP	EUR/JPY	EUR/USD	GBP/USD	USD/CAD	USD/CHF	USD/JPY
Panel (a): Sensitivity to changes in common liquidity									
Whole sample									
a_j	0.5250	−0.1017	−0.0765	−0.0218	−0.1775	0.4448	−0.2451	−0.1949	−0.1990
	(0.0421)	(0.0043)	(0.0164)	(0.0052)	(0.0024)	(0.0160)	(0.0216)	(0.0041)	(0.0032)
b_j	3.1451	0.3541	1.0887	0.5628	0.1725	1.8609	1.6242	0.3483	0.3072
	(0.0655)	(0.0054)	(0.0225)	(0.0066)	(0.0030)	(0.0223)	(0.0322)	(0.0053)	(0.0041)
R^2	0.7594	0.8533	0.7615	0.9075	0.8162	0.9052	0.7765	0.8570	0.8853
Pre-Lehman									
a_j	0.3006	−0.0861	−0.1718	−0.0733	−0.2854	0.2971	0.0716	−0.1370	−0.1264
	(0.0419)	(0.0063)	(0.0232)	(0.0054)	(0.0037)	(0.0267)	(0.0437)	(0.0079)	(0.0048)
b_j	2.7714	0.4106	0.8069	0.4483	−0.0129	1.5113	2.2758	0.4604	0.4263
	(0.0874)	(0.0110)	(0.0434)	(0.0096)	(0.0064)	(0.0492)	(0.0908)	(0.0142)	(0.0085)
R^2	0.7081	0.7706	0.4545	0.8412	0.0095	0.6943	0.6021	0.7175	0.8589
Post-Lehman									
a_j	0.9756	0.0027	−0.3727	−0.0900	−0.1267	0.2224	−0.3639	−0.1842	−0.2589
	(0.1161)	(0.0085)	(0.0348)	(0.0127)	(0.0034)	(0.0345)	(0.0502)	(0.0095)	(0.0069)
b_j	3.6417	0.4405	0.8304	0.5081	0.2204	1.6637	1.4913	0.3542	0.2533
	(0.1434)	(0.0085)	(0.0377)	(0.0128)	(0.0034)	(0.0383)	(0.0585)	(0.0095)	(0.0069)
R^2	0.6724	0.8961	0.6068	0.8331	0.9319	0.8576	0.6743	0.8145	0.8100
Panel (b): Standard deviation of idiosyncratic liquidity									
Whole sample	0.3742	0.0478	0.1245	0.0576	0.0280	0.1004	0.1762	0.0410	0.0384
Pre-Lehman	0.1324	0.0238	0.0716	0.0202	0.0217	0.0697	0.1432	0.0309	0.0243
Post-Lehman	0.5267	0.0665	0.1366	0.0740	0.0345	0.1193	0.2115	0.0507	0.0450

Notes: For each exchange rate j, daily individual liquidity (effective cost), $L_{j,t}^{(ec)}$, is regressed on average FX market-wide liquidity $L_{M,t}^{(ec)}$ (Equation (6)). Liquidity of FX rate j is excluded before computing $L_{M,t}^{(ec)}$. Panel (a) shows the regression results. Heteroscedasticity and autocorrelation (HAC) robust standard errors are shown in parenthesis. Panel (b) shows the standard deviation of idiosyncratic liquidity, which is defined as the residuals of the regression in Equation (6). The sample is January 2, 2007 – December 30, 2009.

Table 8: Descriptive statistics for carry trade returns

Currency	AUD	CAD	DKK	EUR	JPY	NZD	SEK	CHF	GBP
Panel (a): Whole sample									
FX return: $\Delta p_{j,t+1}$									
Mean	−3.58	−3.30	−2.43	−3.43	−8.61	−0.77	−0.21	−5.32	6.34
Std. dev.	20.48	13.93	11.51	11.40	12.93	19.74	16.56	12.15	12.73
Interest rate differential: $i_t^f - i_t^d$									
Mean	2.89	0.06	0.95	0.24	−2.19	3.74	0.26	−1.18	1.09
Std. dev.	1.41	0.73	1.60	1.24	1.93	1.42	1.51	1.28	0.99
Carry trade return: $r_{j,t+1}^e$									
Mean	6.41	3.37	3.36	3.66	6.47	4.44	0.47	4.16	−5.27
Std. dev.	20.47	13.93	11.51	11.39	12.93	19.73	16.56	12.15	12.72
Panel (b): Prior to default of Lehman Brothers									
FX return: $\Delta p_{j,t+1}$									
Mean	−3.01	−6.64	−5.05	−5.11	−6.56	3.09	−1.82	−5.18	2.84
Std. dev.	12.83	9.63	7.87	7.87	10.60	14.39	9.57	9.38	7.95
Interest rate differential: $i_t^f - i_t^d$									
Mean	2.52	−0.14	0.16	−0.10	−3.57	4.01	−0.17	−1.92	1.31
Std. dev.	1.68	0.75	1.44	1.42	1.26	1.46	1.73	1.24	1.06
Carry trade return: $r_{j,t+1}^e$									
Mean	5.48	6.50	5.20	5.01	3.05	0.84	1.65	3.29	−1.55
Std. dev.	12.82	9.63	7.86	7.87	10.60	14.38	9.57	9.37	7.94
Panel (c): After default of Lehman Brothers									
FX return: $\Delta p_{j,t+1}$									
Mean	−4.34	1.10	1.01	−1.21	−11.33	−5.88	1.91	−5.51	10.95
Std. dev.	27.51	18.12	15.04	14.83	15.50	25.14	22.73	15.07	17.11
Interest rate differential: $i_t^f - i_t^d$									
Mean	3.37	0.34	1.98	0.69	−0.37	3.37	0.83	−0.20	0.79
Std. dev.	0.67	0.59	1.16	0.74	0.83	1.27	0.89	0.29	0.81
Carry trade return: $r_{j,t+1}^e$									
Mean	7.65	−0.77	0.93	1.89	10.97	9.18	−1.10	5.31	−10.17
Std. dev.	27.51	18.12	15.04	14.83	15.50	25.13	22.72	15.07	17.10

Notes: This table reports descriptive statistics for different exchange rates with USD being the base currency. Namely, the average log-return, the average interest rate differential as well as daily excess log-returns over UIP are shown. Panel (a) gives results for the whole sample which ranges from January 2, 2007 to December 30, 2009. Summary statistics for two subsamples prior to and after the default of Lehman Brothers are reported in Panels (b) and (c), respectively.

Table 9: Correlation between FX liquidity and carry trade returns

Currency	AUD	CAD	DKK	EUR	JPY	NZD	SEK	CHF	GBP
Panel (a): Daily data									
Liquidity level	0.0800	0.0936	0.0515	0.0529	−0.0797	0.0739	0.0760	0.0135	0.0963
Shocks	0.1907	0.1487	0.1077	0.1154	−0.1391	0.1535	0.1261	0.0295	0.0990
Unexpected shocks	0.1941	0.1577	0.0955	0.1073	−0.1440	0.1549	0.0965	0.0221	0.0819
Panel (b): Monthly data									
Liquidity level	0.3160	0.3562	0.1682	0.1639	−0.2956	0.3318	0.2869	0.0433	0.4785
Shocks	0.6964	0.5953	0.4469	0.4127	−0.4286	0.4986	0.5163	0.1239	0.5263
Unexpected shocks	0.6261	0.5535	0.4357	0.4067	−0.3361	0.4363	0.4664	0.0669	0.3931

Notes: Correlation between FX rate liquidity (level, shocks, unexpected shocks) and carry trade returns with USD being the base currency. Liquidity level is the latent liquidity across exchange rates and liquidity measures extracted by Principal Component Analysis. Shocks and unexpected shocks are obtained as the residuals of an AR(1) and AR(2) model fitted on the liquidity level, respectively. The sample is January 2, 2007 – December 30, 2009.

Table 10: Factor model time series regression results

	AUD	CAD	DKK	EUR	JPY	NZD	SEK	CHF	GBP
				Panel (a): Whole sample					
α	0.0135	0.0059	0.0000	0.0014	0.0180	0.0042	−0.0146	0.0027	−0.0310
	(0.0157)	(0.0174)	(0.0079)	(0.0083)	(0.0157)	(0.0205)	(0.0183)	(0.0126)	(0.0193)
β_{AER}	1.0493	0.6514	1.1083	1.0930	0.6076	1.1565	1.3664	1.1368	0.8308
	(0.0260)	(0.0288)	(0.0132)	(0.0137)	(0.0260)	(0.0340)	(0.0305)	(0.0209)	(0.0321)
β_{IML}	0.3302	0.1968	−0.0899	−0.0912	−0.3818	0.2300	−0.0258	−0.2001	0.0318
	(0.0086)	(0.0095)	(0.0044)	(0.0045)	(0.0086)	(0.0112)	(0.0101)	(0.0069)	(0.0106)
R^2	0.8923	0.7142	0.9127	0.9029	0.7301	0.8021	0.7739	0.8028	0.5760
				Panel (b): Prior to Lehman default					
α	0.0021	0.0145	0.0078	0.0070	0.0096	−0.0171	−0.0085	0.0010	−0.0165
	(0.0153)	(0.0187)	(0.0075)	(0.0076)	(0.0164)	(0.0236)	(0.0146)	(0.0100)	(0.0177)
β_{AER}	1.1698	0.6049	1.0918	1.0921	0.6827	1.2024	1.1976	1.2082	0.7505
	(0.0354)	(0.0434)	(0.0175)	(0.0176)	(0.0381)	(0.0547)	(0.0340)	(0.0232)	(0.0412)
β_{IML}	0.2879	0.2256	−0.0817	−0.0815	−0.4050	0.2980	−0.0363	−0.2328	0.0259
	(0.0111)	(0.0136)	(0.0055)	(0.0055)	(0.0119)	(0.0171)	(0.0106)	(0.0073)	(0.0129)
R^2	0.8523	0.6064	0.9042	0.9035	0.7504	0.7198	0.7562	0.8815	0.4809
				Panel (c): After Lehman default					
α	0.0287	−0.0060	−0.0104	−0.0064	0.0290	0.0291	−0.0202	0.0053	−0.0491
	(0.0299)	(0.0318)	(0.0155)	(0.0164)	(0.0292)	(0.0352)	(0.0377)	(0.0259)	(0.0383)
β_{AER}	0.9815	0.6838	1.1188	1.0983	0.5670	1.1727	1.4261	1.0929	0.8589
	(0.0394)	(0.0420)	(0.0205)	(0.0217)	(0.0385)	(0.0464)	(0.0497)	(0.0341)	(0.0505)
β_{IML}	0.3550	0.1814	−0.0945	−0.0957	−0.3679	0.2010	−0.0282	−0.1820	0.0309
	(0.0133)	(0.0142)	(0.0069)	(0.0073)	(0.0130)	(0.0156)	(0.0168)	(0.0115)	(0.0170)
R^2	0.9072	0.7570	0.9163	0.9032	0.7211	0.8459	0.7835	0.7681	0.6056

Notes: Time series regression results for the daily factor model in Equation (9). $\beta_{AER,j}$ is the factor loading of the market risk factor defined as the average excess FX rate return from the perspective of a US investor. $\beta_{IML,j}$ is the factor loading of the liquidity risk factor defined as the excess return of a portfolio which is long the two most illiquid and short the two most liquid exchange rates. Heteroscedasticity and autocorrelation (HAC) robust standard errors are shown in parenthesis. Panel (a) shows regression results for the whole sample which ranges from January 2, 2007 to December 30, 2009. Regression results for two subsamples prior to and after the default of Lehman Brothers are reported in Panels (b) and (c), respectively.

References

Acharya, V. V., and L. H. Pedersen, 2005, "Asset Pricing with Liquidity Risk", *Journal of Financial Economics*, 77(2), 375–410.

Aït-Sahalia, Y., P. A. Mykland, and L. Zhang, 2005, "How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise", *The Review of Financial Studies*, 18(2), 351–416.

Amihud, Y., 2002, "Illiquidity and Stock Returns: Cross-Section and Time-Series Effects", *Journal of Financial Markets*, 5(1), 31–56.

Amihud, Y., and H. Mendelson, 1986, "Asset Pricing and the Bid-Ask Spread", *Journal of Financial Economics*, 17(2), 223–249.

Bank for International Settlements, 2010, "Foreign Exchange and Derivatives Market Activity in April 2010", Triennial Central Bank Survey.

Berger, D. W., A. P. Chaboud, S. V. Chernenko, E. Howorka, and J. H. Wright, 2008, "Order Flow and Exchange Rate Dynamics in Electronic Brokerage System Data", *Journal of International Economics*, 75, 93–109.

Bernanke, B. S., 2008, "Risk Management in Financial Institutions", Speech at the Federal Reserve Bank of Chicago's Annual Conference on Bank Structure and Competition, Chicago, Illinois, May 15, 2008.

Bjønnes, G. H., C. L. Osler, and D. Rime, 2008, "Asymmetric Information in the Interbank Foreign Exchange Market", Working paper 2008-25, Central Bank of Norway, Research Department.

Breedon, F., and P. Vitale, 2010, "An Empirical Study of Portfolio-Balance and Information Effects of Order Flow on Exchange Rates", *Journal of International Money and Finance*, 29(3), 504–524.

Brownlees, C. T., and G. M. Gallo, 2006, "Financial Econometric Analysis at Ultra-High Frequency: Data Handling Concerns", *Computational Statistics & Data Analysis*, 51(4), 2232–2245.

Brunnermeier, M. K., S. Nagel, and L. H. Pedersen, 2009, "Carry Trades and Currency Crashes", *NBER Macroeconomics Annual 2008*, 23, 313–347.

Brunnermeier, M. K., and L. H. Pedersen, 2009, "Market Liquidity and Fund Liquidity", *The Review of Financial Studies*, 22(6), 2201–2238.

Burnside, C., 2009, "Carry Trades and Currency Crashes: A Comment", *NBER Macroeconomics Annual 2008*.

Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo, 2011, "Do Peso Problems Explain the Returns to the Carry Trade?", *The Review of Financial Studies*, 24(3), 853–891.

Campbell, J. Y., S. J. Grossman, and J. Wang, 1993, "Trading Volume and Serial Correlation in Stock Returns", *The Quarterly Journal of Economics*, 108(4), 905–939.

Chaboud, A. P., S. V. Chernenko, and J. H. Wright, 2007, "Trading Activity and Exchange Rates in High-Frequency EBS Data", International Finance Discussion Papers 903, Board of Governors of the Federal Reserve System and Harvard University.

Chaboud, A. P., B. Chiquoine, E. Hjalmarsson, and C. Vega, 2009, "Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market", International Finance Discussion Papers 980, Washington: Board of Governors of the Federal Reserve System.

Chordia, T., R. Roll, and A. Subrahmanyam, 2000, "Commonality in Liquidity", *Journal of Financial Economics*, 56(1), 3–28.

———, 2001, "Market Liquidity and Trading Activity", *Journal of Finance*, 56(2), 501–530.

Engel, C., 1992, "The Risk Premium and the Liquidity Premium in Foreign Exchange Markets", *International Economic Review*, 33(4), 871–879.

Evans, M. D. D., and R. K. Lyons, 2002, "Order Flow and Exchange Rate Dynamics", *Journal of Political Economy*, 110(1), 170–180.

Fama, E. F., 1984, "Forward and Spot Exchange Rates", *Journal of Monetary Economics*, 14(3), 319–338.

Fama, E. F., and K. R. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds", *Journal of Financial Economics*, 33(1), 3–56.

Farhi, E., S. P. Fraiberger, X. Gabaix, R. Ranciere, and A. Verdelhan, 2009, "Crash Risk in Currency Markets", NBER Working paper 15062, Harvard University, New York University, International Monetary Fund, and Boston University.

Foucault, T., O. Kadan, and E. Kandel, 2005, "Limit Order Book as a Market for Liquidity", *The Review of Financial Studies*, 18(4), 1171–1217.

Gilmore, S., and F. Hayashi, 2008, "Emerging Market Currency Excess Returns", NBER Working paper 14528, AIG Financial Products and University of Tokyo.

Goyenko, R. Y., C. Holden, and C. Trzcinka, 2009, "Do Liquidity Measures Measure Liquidity?", *Journal of Financial Economics*, 92(2), 153–181.

Hampel, F. R., E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel, 2005, *Robust Statistics : The Approach Based on Influence Functions*. Wiley, New York, USA.

Hansen, L. P., and R. J. Hodrick, 1980, "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis", *The Journal of Political Economy*, 88(5), 829.

Hasbrouck, J., 2009, "Trading Costs and Returns for US Equities: Estimating Effective Costs from Daily Data", *The Journal of Finance*, 64(3), 1445–1477.

Hasbrouck, J., and D. Seppi, 2001, "Common Factors in Prices, Order Flows, and Liquidity", *Journal of Financial Economics*, 59(3), 383–411.

Hodrick, R., and S. Srivastava, 1986, "The Covariation of Risk Premiums and Expected Future Spot Rates", *Journal of International Money and Finance*, 3, 5–30.

Holden, C. W., 2009, "New Low-Frequency Spread Measures", *Journal of Financial Markets*, 12, 778–813.

Jones, C. M., G. Kaul, and M. L. Lipson, 1994, "Transactions, Volume, and Volatility", *The Review of Financial Studies*, 7(4), 631–651.

Jurek, J. W., 2009, "Crash-Neutral Currency Carry Trades", Working paper, Princeton University.

Korajczyk, R. A., and R. Sadka, 2008, "Pricing the Commonality Accross Alternative Measures of Liquidity", *Journal of Financial Economics*, 87(1), 45–72.

Kyle, A., 1985, "Continuous Auctions and Insider Trading", *Econometrica*, 53, 1315–1335.

Lee, C., and M. Ready, 1991, "Inferring Trade Direction from Intraday Data", *Journal of Finance*, 46(2), 733–746.

Liu, W., and A. Maynard, 2005, "Testing Forward Rate Unbiasedness Allowing for Persistent Regressors", *Journal of Empirical Finance*, 12(5), 613–628.

Lustig, H. N., N. L. Roussanov, and A. Verdelhan, 2010, "Common Risk Factors in Currency Markets", Working paper, UCLA, University of Pennsylvania, and MIT.

Maronna, R. A., R. D. Martin, and V. J. Yohai, 2006, *Robust Statistics: Theory and Methods*. John Wiley & Sons, Chichester, West Sussex, UK.

Marsh, I. W., and C. O'Rourke, 2005, "Customer Order Flow and Exchange Rate Movements: Is there Really Information Content?", Working paper, Cass Business School.

Melvin, M., and M. P. Taylor, 2009, "The Crisis in the Foreign Exchange Market", *Journal of International Money and Finance*, 28(8), 1317–1330.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf, 2011, "Carry Trades and Global Foreign Exchange Volatility", *The Journal of Finance*, forthcoming.

Morris, S., and H. S. Shin, 2004, "Liquidity Black Holes", *Review of Finance*, 8(1), 1–18.

Pástor, L., and R. F. Stambaugh, 2003, "Liquidity Risk and Expected Stock Returns", *Journal of Political Economy*, 111(3), 642–685.

Perold, A. F., 1988, "The Implementation Shortfall: Paper vs. Reality", *Journal of Portfolio Management*, 14(3), 4–9.

Plantin, G., and H. S. Shin, 2010, "Carry Trades, Monetary Policy and Speculative Dynamics", Working paper, London Business School and Princeton University.

Ranaldo, A., and P. Söderlind, 2010, "Safe Haven Currencies", *Review of Finance*, 14(3), 385–407.

Roll, R., 1984, "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market", *Journal of Finance*, 39(4), 1127–1139.

Ross, S. A., 1976, "The Arbitrage Theory of Capital Asset Pricing", *Journal of Economic Theory*, 13, 341–360.

Rosu, I., 2009, "A Dynamic Model of the Limit Order Book", *The Review of Financial Studies*, 22(11), 4601–4641.

Sadka, R., 2006, "Momentum and Post-earnings Announcement Drift Anomalies", *Journal of Financial Economics*, 80, 309–349.

World Federation of Exchanges, 2009, *2009 Annual Report and Statistics*. Paris, France.

Zhang, L., P. A. Mykland, and Y. Aït-Sahalia, 2005, "A Tale of Two Time Scales: Determining Integrated Volatility With Noisy High-Frequency Data", *Journal of the American Statistical Association*, 100(472), 1394–1411.

The Joint Dynamics of Hedge Fund Returns, Illquidity, and Volatility

Jan Wrampelmeyer

This paper was presented at:

- LSE Alternative Investment Research Conference, London, UK (March 2011)

The most recent version of the paper is available at `http://ssrn.com/abstract=1720909`

Abstract

Hedge funds are frequently blamed for increasing volatility and illiquidity in financial markets. I investigate the validity of this hypothesis by modeling the joint dynamics of hedge fund returns and volatility as well as illiquidity in the equity and the foreign exchange (FX) market. The results show that hedge funds tend to profit from periods of low equity liquidity, but react negatively to shocks in volatility and FX illiquidity, indicating a significant FX exposure of many strategies. Moreover, I find weak evidence that hedge funds cause higher volatility in financial markets with trend following strategies being the main transmission channel. Finally, there exist cross-market dynamics and bidirectional spillovers between volatility and illiquidity in the equity and FX market. These results have important implication for performance attribution, risk management as well as regulatory policy.

1. Introduction

During the last two decades the size and importance of the hedge fund industry have increased tremendously with significant consequences for financial markets. According to estimates from International Financial Services (2010), assets under management by hedge funds have grown from \$50*bn* in 1990 to over \$2,100*bn* in 2007 and are still estimated to total \$1,700*bn* at the end of 2009. In the first quarter of 2010, hedge fund related trading accounted for almost one third of the average US daily share volume (Harris, 2010).

Being part of the shadow banking system, hedge funds are largely unregulated with no lender of last resort, no discount window, no deposit insurance and no capital requirement regulation (Sullivan, 2008). Therefore, in light of the recent financial crisis, the role of hedge funds in financial markets is controversial and gave rise to an increasing level of scrutiny by politicians and regulators. On the one hand, hedge funds are acknowledged to be important liquidity providers which increase market efficiency and help to hold corporate management accountable. On the other hand, hedge funds have recently been blamed for high volatility as well as liquidity problems and market crashes. For instance in summer 2009, Lord Turner, chairman of the British Financial Services Authority, described much of the City of London's activities as "socially useless" (Turner, 2009) and later added at the World Economic Forum in Davos that hedge funds in the form of carry traders add no value, but rather pose a risk to the real economy. More recently, the German government instituted a ban on naked short sales of the ten most important financial stocks as well as credit default swaps (CDS) on Euro government bonds, because politicians accused speculators to be responsible for the strong decline of the Euro during the Greek debt crisis and for the elevated volatility of financial stocks during that period.

There exist a number of potential explanations for this public concern: Hedge funds are typically organized as private investment vehicles for institutional investors and high net worth individuals. Moreover, they tend to be secretive and do not disclose their activities to the public. This non-traditional business model evokes uncertainty about market quality and fear that volatility and illiquidity are magnified by hedge fund activity. Furthermore, the recent run for speed among statistical arbitrage funds violates the common opinion that long term investing is superior to short term speculation. Finally, the ability to take on leverage and short positions leads to more complex risk profiles compared to traditional, long-only funds.

The goal of this paper is to shed light on the relation between hedge fund returns, volatility and illiquidity in the equity as well as foreign exchange market. A thorough understanding of these empirical linkages is vital given that new regulation, such as, e.g., more stringent capital requirements for banks, might lead to further growth of the shadow banking system. Analyzing the joint dynamics facilitates the detection of, potentially bidirectional, causalities between the market variables and hedge fund returns, which are of importance not only for investors and hedge fund managers, but also for regulators as well as politicians: For investors and fund of funds, a proper understanding of hedge funds' risk dynamics is crucial for performance attribution and evaluation (Bollen and Whaley, 2009). Moreover, it supports fund of funds' decisions to allocate assets to different managers and to diversify effectively. Apart from investing in stocks, hedge funds trade extensively in foreign exchange markets, for instance, using the famous carry trade. Therefore managers profit from evaluating their exposure to volatility and illiquidity in the two markets as this analysis supports portfolio allocation and hedging decisions. Moreover, a joint modeling of the dynamics allows to quantify potential responses of market variables to shocks in hedge fund returns, helping regulators and politicians to assess the effect of hedge fund activity on volatility and illiquidity. Analyzing volatility and illiquidity in both the equity as well as the FX market enables market participants to detect feedback effects and spillovers between the two markets, which have important implications, for instance in the context of portfolio selection and risk management.

The results show that various hedge fund strategies respond negatively to shocks in equity and FX volatility. Thus, hedge funds do not profit from higher levels of volatility. Similar results are obtained for FX illiquidity, indicating a significant FX exposure of a number of hedge fund strategies that is not captured by existing factor models for hedge fund returns. In contrast, I find evidence that hedge fund returns are larger following periods of equity illiquidity supporting the hypothesis that hedge funds are able to profit from trading opportunities during illiquid market conditions. Moreover, there is some evidence that hedge funds relying on computerized trading and trend following strategies cause higher volatility in financial markets. On the other hand dedicated short bias fund returns actually decrease volatility and there is no indication that hedge fund returns increase illiquidity. Finally, the paper documents cross-market dynamics and spillover effects between volatility and illiquidity in the equity as well as the FX market.

The remainder of this paper is structured as follows: The first section presents the hypotheses and briefly reviews the relevant literature. In the second section, the data as well as measures for volatility and illiquidity are introduced. The third section models the joint dynamics of hedge fund returns, volatility and illiquidity to analyze the hypotheses and to provide evidence for potential bidirectional causalities. The final section provides a summary and concluding remarks.

2. Hypotheses and Review of the Related Literature

Public concern and increasing scrutiny about a potentially destabilizing role of hedge funds in financial markets in general, and the interrelation between hedge fund returns, volatility, and market illiquidity in particular, can be summarized by the following hypotheses that will be investigated in this paper:

(i) Illiquidity and volatility in both the FX and equity market positively impact hedge fund returns.

(ii) Hedge fund activity increases illiquidity and volatility in both the FX and equity market.

(iii) There are illiquidity and volatility spillovers between the FX and stock market.

The first hypothesis relates to the question whether hedge funds profit from periods of high volatility and illiquidity. Alternatively, these periods might pose a risk to hedge funds leading to lower or even negative returns. There exist various reasons to believe that hedge fund returns are influenced by volatility in equity markets. Many hedge fund strategies invest in stocks and are, thus, subject to equity volatility risk. Next to strategies like long/short equity or dedicated short bias, in particular the event-driven investment approach has significant exposure to short volatility risk (Anson and Ho, 2003). On the other hand, equity market neutral funds might profit from an increase in volatility by applying gamma trading strategies. Hence, volatility in the market might both be dangerous as well as beneficial for hedge funds.

Similarly, changes in equity liquidity might severely impact hedge fund returns, because many hedge funds apply leverage and invest in complex and potentially illiquid securities. Cao, Chen, Liang, and Lo (2010) show that hedge funds adjust their market exposure in light of changing market liquidity, which potentially also affects their returns. Moreover, Sadka (2010) documents that unexpected changes in stock market liquidity are a priced risk factor in the cross section of hedge fund returns. Finally, Boyson, Stahel, and Stulz (2010) find that large declines in stock market liquidity can cause contagion in hedge fund returns.

Not all hedge funds invest (solely) in domestic equity markets. A frequently applied strategy, for instance, by global macro hedge funds is the carry trade. Mancini, Ranaldo, and Wrampelmeyer (2010) as well as Menkhoff, Sarno, Schmeling, and Schrimpf (2011) show that both FX liquidity as well as FX volatility are risk factors in international asset pricing models helping to explain carry trade returns. Moreover, US based funds which invest abroad are exposed to currency movements. Thus, trading foreign exchange is attractive not only to speculators, but also for hedging and risk management purposes in conjunction with foreign equities portfolio holdings (Campbell, Serfaty-de Medeiros, and Viceira, 2010). Due to these strategies, movements in FX volatility and illiquidity might significantly impact hedge fund returns.

The second hypothesis reverses the causality between hedge fund returns and illiquidity as well as volatility. It is highly controversial whether hedge funds influence illiquidity and volatility in financial markets. On the one hand, hedge funds are seen as liquidity providers and sophisticated investors who increase market efficiency. Aragon and Strahan (2009) show that market liquidity of stocks traded by hedge funds with Lehman Brothers as prime broker declined stronger than other stocks following the default of Lehman in September 2008. Thus, a decrease in hedge funds' funding liquidity has a negative effect on market liquidity, supporting the argument that hedge funds provide liquidity. Hedge fund managers constantly scan the market for trading opportunities and are ready to trade on these beliefs, moving prices towards equilibrium. In particular the use of short selling is beneficial, because overpriced securities are driven towards their fundamental value. This reduces volatility and facilitates the price discovery process leading to fewer asset price bubbles and better allocation of resources in the real economy. In line with this reasoning, Hendershott, Jones, and Menkveld (2010) find that the increasing volume of algorithmic trading, which is the means of trading of statistical arbitrage hedge funds, improves liquidity and enhances the informativeness of quotes. According to this theory, positive hedge fund returns should imply lower volatility and illiquidity. On the other hand, some people view hedge funds as speculators which take excessive risks and earn money by driving prices away from their fundamental value, therefore, increasing volatility and causing illiquidity, for instance, through crowded trades in one direction. Furthermore, even if hedge funds make markets more efficient, this might actually increase volatility, because market participants have the tendency to react in the same way to unanticipated news leading to larger price movements. Based on these arguments, the German government recently instituted a ban on naked short sales in an effort to limit volatility. This theory implies a positive relation between hedge fund returns and illiquidity as well as volatility.

The final hypothesis investigates the linkage between illiquidity as well as volatility in equity and foreign exchange markets. This is related to the study of Chordia, Sarkar, and Subrahmanyam (2005) which establishes that common factors drive bond and equity liquidity. Moreover, Goyenko and Ukhov (2009) document a lead/lag relation between illiquidity in these markets. Fleming, Kirby, and Ostdiek (1998) find similar linkages for volatility. Given the tremendous size and importance of the FX market[1], this paper analyzes spillovers between illiquidity in equity and FX markets. There exists a number of reasons to expect a connection between equity and FX illiquidity: If liquidity dries up in FX market, which is the world's largest financial market, this is a good indication for a liquidity crisis with effects in all financial markets. Mancini, Ranaldo, and Wrampelmeyer (2010) show that FX liquidity decreased significantly during 2007–2009 and that FX liquidity is correlated to measures of equity liquidity. This interdependence is consistent with the interaction of market and funding liquidity during liquidity spirals (Brunnermeier and Pedersen, 2009). Finally, central bank interventions directly impact the FX market, but have severe implications for other markets and the worldwide economy, implying a potential lead/lag relation between illiquidity and volatility. For instance, if monetary shocks are reflected in FX market illiquidity first, this might influence future equity illiquidity. The following section describes the data set that is used to explore all of these hypotheses.

3. The Data

3.1. Hedge Fund Indexes and Benchmark Model for Hedge Fund Returns

This paper relies on the Dow Jones Credit Suisse Hedge Fund Index[2] to measure hedge fund returns. With data going back to 1994, the index is calculated as well as rebalanced monthly and considers all funds included in the Lipper TASS database that (i) have at least USD 50 million assets under management, (ii) provide current audited financial statements, and (iii) have a at least a one year track record. The funds in the index are asset weighted implying that the index accurately depicts the performance of the hedge fund industry. Furthermore, Dow Jones Credit Suisse provides subindexes corresponding to various hedge fund styles[3], accommodating for the

[1]With an estimated average daily turnover of four trillion US dollar in 2010 (Bank for International Settlements, 2010), the FX market is by far the world's largest financial market. This volume corresponds to more than ten times that of global equity markets (World Federation of Exchanges, 2009).

[2]This index was formerly known as CSFB/Tremont as well as Credit Suisse/Tremont.

[3]The internet appendix, which can be downloaded from `http://www.wrampelmeyer.com`, contains a description of the different hedge fund styles.

diversity within the hedge fund industry.

An ubiquitous problem with hedge fund data is caused by the funds' opaqueness and the lack of uniform reporting standards. This potentially results in a number of biases when investigating hedge fund returns. In an effort to minimize the most severe data problems, namely the backfilling and survivorship bias, hedge funds joining the Dow Jones Credit Suisse Hedge Fund index are not allowed to backfill their historical returns. Thus, there is no backfill or instant history bias resulting from the construction of the index. Moreover, survivorship bias is minimized because funds in the process of liquidation are not removed from the index. Therefore, the index captures the potential negative performance before a fund ceases to exist. As the purpose of this paper is to investigate the dynamics of hedge fund returns, rather than to compare absolute performance measures, relying on hedge fund indexes is a viable option as they are useful for describing the risk characteristics of hedge fund strategies (Jaeger and Wagner, 2005).

While the goal of many hedge funds is to produce absolute returns, various recent studies document evidence for hedge fund betas. In a series of papers, Fung and Hsieh (1997a, 2001, 2004) extend classic linear factor models to explain hedge fund returns. Jaeger and Wagner (2005) as well as Hasanhodzic and Lo (2007) describe hedge fund replication strategies based on a number of traditional and non-traditional risk factors. Additional support for the presence of systematic risk factors in hedge fund returns became apparent during the quant crisis of 2007 during which almost all long/short equity hedge funds suffered severe losses (Khandani and Lo, [2007, 2011]).

To control for the exposure to common risk factors, the hedge fund index returns are filtered using the Fung-Hsieh seven factor model. First, this frequently applied model includes standard equity risk factors such as the excess market return ($MKT - RF$) and the spread between small-cap and large-cap stock returns (SMB) of Fama and French (1993). Second, fixed income exposure is covered by changes in the 10-year treasury yield ($\Delta TERM$) and changes in the credit spread ($\Delta CREDIT$ = Moody's Baa yield minus 10-year treasury yield)[4]. Finally, the model contains trend following factors based on portfolios of lookback straddles on bonds ($PFTSBD$), exchange rates ($PTFSFX$), and commodities ($PTFSCOM$) to capture option like features in hedge fund returns (see also Agarwal and Naik (2004)) and the risk inherent in dynamic investment styles[5]. Using excess returns prevents that effects from known common risk factors influence the relation between hedge fund returns, illiquidity, and volatility. In this paper, excess returns are denoted by $r_{HF,all}$, where the second subscript refers to the specific hedge fund style index. Thus, $r_{HF,all}$ is the excess return of the Dow Jones Credit Suisse Hedge Fund Index for all hedge funds, whereas, for instance, $r_{HF,em}$ denotes emerging markets hedge funds' excess returns.

---

[4]Data on the 10-year treasury constant maturity yield and Moody's Baa yield are obtained from the website of the Federal Reserve `http://www.federalreserve.gov/econresdata/releases/statisticsdata.htm`.

[5]I thank David Hsieh and Kenneth French for providing their risk factors on their respective websites: `http://faculty.fuqua.duke.edu/\%7Edah7/DataLibrary/TF-Fac.xls` and `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`.

3.2. Measures of Equity Volatility and Illiquidity

Estimates for monthly equity volatility and illiquidity are obtained from Center for Research in Security Prices (CRSP) data for all listed NYSE, AMEX, and NASDAQ common stocks. This comprehensive database of accurate stock prices ensures that the proxies adequately represent the US equity market.

Amihud's (2002) measure is used as proxy for illiquidity in the stock market. The measure is based on daily data and captures the price impact of order flow which is an important aspect of illiquidity (Kyle, 1985).[6] More precisely, illiquidity of stock i in month t is defined as

$$ILL^i_{EQ,t} = \frac{1}{D^i_t} \sum_{d=1}^{D^i_t} \frac{\left| r^i_{t,d} \right|}{v^i_{t,d}}, \tag{1}$$

where D^i_t is the number of valid observations[7] and $r^i_{t,d}$ and $v^i_{t,d}$ are the return and dollar volume (in millions) on day d in month t. The intuition of this illiquidity measure is that if a stock is illiquid, trading volume will be associated with a strong effect on the stock price. Thus, an illiquid stock will exhibit a large value of $ILL^i_{EQ,t}$. Given estimates for the illiquidity of individual stocks, the illiquidity of the equity market $ILL_{EQ,t}$ is computed subsequently by averaging $ILL^i_{EQ,t}$ over all valid stocks[8].

Equity volatility in month t is proxied straightforwardly by averaging daily absolute returns of the value weighted CRSP market index, $r^m_{t,d}$:

$$VOL_{EQ,t} = \frac{1}{D_t} \sum_{d=1}^{D^i_t} \left| r^m_{t,d} \right|, \tag{2}$$

where D_t denotes the number of trading days in month t. The use of absolute returns allows for better comparison to the measure of foreign exchange volatility which is introduced next.

3.3. Measures of Foreign Exchange Volatility and Illiquidity

Daily data for USD based spot exchange rates are obtained from Reuters (via Datastream) for the following 47 countries: Australia, Austria, Belgium, Brazil, Bulgaria, Canada, Croatia, Cyprus, Czech Republic, Denmark, Egypt, Finland, France, Germany, Greece, Hong Kong, Hungary, Iceland, India, Indonesia, Ireland, Israel, Italy, Japan, Kuwait, Malaysia, Mexico, The Netherlands, New Zealand, Norway, Philippines, Poland, Portugal, Russia, Saudi Arabia, Singapore, Slovakia, Slovenia, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan, Thailand, Ukraine, and the United Kingdom. Since the institution of the European Monetary Union and the introduction of the Euro in January 1999, the sample is successively reduced to

[6] In a comparison of various liquidity proxies, Goyenko, Holden, and Trzcinka (2009) show that Amihud's (2002) measure is a good proxy for the price impact of trades.

[7] A day is a valid observation for stock i if the stock is traded during the day, i.e., if the dollar volume is positive.

[8] In line with Goyenko and Ukhov (2009), $ILL^i_{EQ,t}$ is a valid observation in month t if $D^i_t \geq 15$.

34 countries with the Euro area replacing the individual currencies of the member states.

As proxy for FX volatility I rely on a measure similar to the one for equity markets. More precisely, in line with Menkhoff, Sarno, Schmeling, and Schrimpf (2011), global FX volatility in month t is defined as the average absolute return over all exchange rates and trading days in that month:

$$VOL_{FX,t} = \frac{1}{D_t} \sum_{d=1}^{D_t^i} \left[\frac{1}{N_d} \sum_{n=1}^{N_d} \left| r_{t,d}^i \right| \right], \tag{3}$$

with N_d denoting the number of currencies in the cross section on day d. Since, the sample contains several emerging markets, absolute exchange rate returns are used to alleviate the effect of potential outliers.

FX illiquidity is estimated by bid-ask spreads[9] which are also available from Datastream:

$$ILL_{FX,t} = \frac{1}{D_t} \sum_{d=1}^{D_t^i} \left[\frac{1}{N_d} \sum_{n=1}^{N_d} \frac{s_{t,d}^{ask,i} - s_{t,d}^{bid,i}}{0.5(s_{t,d}^{ask,i} + s_{t,d}^{bid,i})} \right]. \tag{4}$$

Note that $s_{t,d}^{bid,i}$ and $s_{t,d}^{ask,i}$ represent the bid and ask quotes of exchange rate i on day d, respectively. Unfortunately, Reuters bid-ask quotes on Datastream are only indicative. For that reason, I test the accuracy of these quoted spreads by comparing the illiquidity proxy for selected exchange rates to the benchmark spread measures presented by Mancini, Ranaldo, and Wrampelmeyer (2010). During 2007–2009, the correlation between transactable benchmark spreads and the proxy based on indicative quotes ranges from a moderate 0.511 for USD/CHF to a high 0.898 for monthly USD/JPY data[10].

Summary statistics for all illiquidity and volatility measures are presented in Table 1. As expected the foreign exchange market is less volatile than the equity market. Furthermore, illiquidity in both markets varied substantially during the sample. In the equity market, the price impact was ten times as high in illiquid periods compared to the most liquid parts of the sample. In the FX market, the bid-ask spreads peaks at 27.16 basis points which is more than 4.5 times the minimum. The next section will investigate the dynamic relation between these market variables and hedge fund returns.

[Table 1 about here.]

4. The Joint Dynamics of Hedge Fund Returns, Illiquidity and Volatility

4.1. Preliminary Analysis

In order to investigate the contemporaneous relation between hedge fund returns, illiquidity and volatility, Table 2 depicts pairwise correlations between these endogenous variables. The

[9]The measure of FX illiquidity differs from the one in equity markets due to the lack of detailed FX volume data. However, the results of Mancini, Ranaldo, and Wrampelmeyer (2010) show that FX liquidity measures covary significantly and exhibit similar time series patterns and cross-sectional variation. Consequently, bid-ask spreads are a reliable and easy to compute measure of FX illiquidity.

[10]Note that correlations for daily data are much lower (0.255 for USD/CHF to 0.635 for USD/JPY).

excess returns of various hedge funds strategies exhibit significant positive correlation indicating common return components additional to the seven Fung-Hsieh factors. Notably, the equity market neutral and dedicated short bias strategies are least linked to the other strategies with correlations even being negative in certain cases.

Volatility in both the equity as well as the FX market is negatively related to excess hedge fund returns, i.e., hedge fund returns tend to be lower in periods of high volatility. More precisely, FX volatility is negatively related to all strategies; the same holds true for equity volatility except for convertible arbitrage funds. The correlations between the excess returns of various strategies and the illiquidity proxies of the two markets are more intricate. Illiquidity in the equity market is significantly positively correlated with convertible arbitrage, equity market neutral, and fixed income arbitrage hedge fund returns. Thus, these hedge funds appear to earn high returns in periods of low equity liquidity. On the other hand, illiquidity in FX markets is contemporaneously negatively correlated to the returns of the event driven, emerging markets, fixed income arbitrage and global macro strategies.

Volatility and illiquidity in the equity and foreign exchange market are contemporaneously correlated. This correlation is particularly strong for the volatility across the two markets. Overall, the correlation structure indicates that volatility and illiquidity in the equity and in the FX market comove, however, their dynamic relation with hedge fund returns might differ not only across market variables, but also across hedge fund styles.

[Table 2 about here.]

4.2. VAR Model

In the following, the potential dynamic effects and bidirectional causalities between equity and FX volatility and illiquidity as well as excess hedge fund returns will be investigated to provide insights about the hypotheses postulated above. To model the joint dynamics of hedge fund returns, illiquidity and volatility, I adopt a vector autoregression (VAR) specification including five endogenous variables:

$$\mathbf{X}_t = c + \sum_{j=1}^{J} \mathbf{A}_j \mathbf{X}_{t-j} + \mathbf{u}_t, \tag{5}$$

where $\mathbf{X}_t = [r_{HF,t}, VOL_{EQ,t}, ILL_{EQ,t}, VOL_{FX,t}, ILL_{FX,t}]'$. The lag length J is selected by Akaike and Schwarz information criteria. Given that hedge funds are extremely heterogeneous and use a variety of diverse strategies, as also confirmed by the correlations in Table 2, pooling all hedge fund returns might result in a significant loss of information. Therefore, it is relevant to not only rely on $r_{HF,t}$, but also investigate the dynamics separately based on different style indexes. Thus, I estimate analogous models with excess returns of individual hedge fund styles replacing the overall hedge fund market index. All endogenous variables and VAR specifications are tested for stationarity finding no evidence that the stationarity condition[11] is violated.

A first test for causal relations between the endogenous variables can be based on Granger's (1969) concept of causality. Intuitively, a variable y_t is said to Granger cause variable z_t if the former helps to improve forecasts for the latter. In the VAR setting, testing for Granger causality is equivalent to testing whether the coefficients of $y_{t-1}, \ldots, y_{t-J}$ are jointly equal to zero in the equation with z_t as the dependent variable. Tables 3, 4, and 5 show χ^2-statistics together with the corresponding p-values for pairwise Granger causality tests between the endogenous variables.

Table 3 provides evidence for the hypothesis that volatility and illiquidity in the equity and foreign exchange market Granger cause hedge fund returns. For almost all hedge fund strategies at least one of the market variable Granger causes hedge fund returns, with evidence being particularly strong for equity market neutral and global macro funds. Interestingly, foreign exchange illiquidity helps to forecast the most hedge fund strategy returns, while equity volatility is only important for the returns of equity market neutral funds.

The second hypothesis is analyzed in Table 4 showing that evidence for bidirectional causalities and hedge fund returns helping to forecast market variables is rather limited. There is some evidence that dedicated short bias and equity market neutral funds Granger cause VOL_{EQ}. Moreover, event driven fund returns help to forecast FX volatility and illiquidity.

Finally, results concerning the third hypothesis are presented in Table 5 indicating that equity volatility and FX illiquidity play the leading role in the cross-market dynamics. While there is evidence that all other variables Granger cause ILL_{EQ} and VOL_{FX}, no variable significantly improves forecasts for ILL_{FX}. Similarly, VOL_{EQ} is not Granger caused by illiquidity in the two markets and the null hypothesis that VOL_{FX} improves forecasts of VOL_{EQ} can not be rejected only at a significance level of 25%.

[Tables 3, 4, and 5 about here.]

A limitation of pairwise Granger causality tests is that they only provide insights with regard to the relation between two variables from a single VAR equation; the joint dynamics implied by the full VAR specification are neglected. Therefore, potentially bidirectional dynamics taking feedback effects into account are investigated using impulse response functions (IRFs) in the next subsection.

4.3. Impulse Response Analysis

Impulse response functions visualize how a single positive shock to one variable affects current and future values of all endogenous variables within the VAR model. For instance, cross-market dynamics in illiquidity can be analyzed by tracing the effect of a shock to FX illiquidity on equity illiquidity over time. Given that the innovations in the VAR model are likely to be correlated,

[11] VAR model (5) is stationary, if all roots of the characteristic polynomial lie outside the unit circle or equivalently det $\left(\mathbf{I} - (\sum_{j=1}^{J} \mathbf{A}_j z^j)\right) \neq 0$ for $|z| \leq 1$. For individual variables I rely on the Augmented DickeyFuller and the Phillips-Perron unit root tests.

they are orthogonalized using a Cholesky decomposition. Due to this orthogonalization, results from impulse response analyses might depend on the ordering of endogenous variables. However, for the model at hand, the internet appendix shows that the conclusions about the interrelations of market variables and hedge fund returns are insensitive to the sequence in which shocks enter the system.[12]

Figures 1, 2, and 3 show IRFs of various endogenous variables to a positive one standard deviation shock to one of the variables. All graphs trace the response of a variable to a unit standard deviation change in another variable over a period of 24 months. To indicate the statistical significance of the responses, the dotted lines depict bootstrapped 90% confidence intervals.

Figure 1 depicts the response of the returns of various hedge fund strategies to shocks in equity and foreign exchange volatility and illiquidity. As shown in Panel (a), the effect of an impulse to equity volatility is negative for the returns of almost all hedge fund strategies. Interestingly, this effect is most pronounced for equity market neutral funds. A potential explanation for the lack of neutrality with respect to shocks to the second moment of equity market returns is that an increase in volatility has a negative impact on the hedging ability of funds within this style category. Moreover, the dedicated short bias and long-short equity strategies are negatively affected, which is in line with the funds' short volatility risk. Albeit not statistically significant, there is some evidence that emerging markets and convertible arbitrage funds profit from an increase in equity volatility. Given that convertible arbitrage funds are typically long the (underpriced) call option embedded in convertible bonds, this positive reaction to a shock in equity volatility is intuitive given the positive vega[13] of the option.

In contrast, Panel (b) shows that almost all hedge fund strategies experience an increase in returns following a shock to equity illiquidity, with the effect being strongest for equity market neutral and global macro funds. This is in line with the argument that hedge fund managers are sophisticated investors that might have profited from trading opportunities during the crisis when liquidity was extremely low and many assets were relatively cheap. The only exception are commodity trading advisors (CTAs) in the managed futures category for which equity illiquidity has a (not significantly) negative influence.

Panel (c) of 1 Appendix A shows impulse response functions of various hedge fund styles to a shock in FX volatility. In general, a more volatile FX market leads to lower returns for all hedge fund styles with managed futures again being the only non-significant exception. For fixed income or global macro funds, the carry trade is a popular investment strategy which is exposed to FX volatility risk (Menkhoff, Sarno, Schmeling, and Schrimpf, 2011). More surprising are the negative responses of equity related strategies such as long/short equity or equity market neutral. A possible explanation for this sensitivity to FX volatility is a significant FX exposure of these hedge funds due to investments in foreign markets even after controlling for the currency

[12]The internet appendix presents results for the following order of variables is: VOL_{FX}, VOL_{EQ}, ILL_{FX}, ILL_{EQ}, r_{HF}.

[13]The vega of an option denotes the derivative of the option value with respect to the volatility of the underlying.

trend following factor $PFTSFX$.

Impulse responses to shocks in FX illiquidity are shown in Panel (d). Emerging market and in particular global macro hedge fund returns react negatively to shocks in FX illiquidity. For the latter, this is in line with Mancini, Ranaldo, and Wrampelmeyer (2010) who find that FX illiquidity poses a significant risk to FX investors in general and carry traders in particular. The negative response of emerging market hedge funds might be related to illiquidity cost when unwinding positions denominated in foreign currency. In contrast, managed future funds react positively to an increase in FX illiquidity. An explanation for this positive response might be that FX managed future managers tend to profit in periods of illiquidity, potentially by being able to exploit arbitrage opportunities[14] in times of distressed market conditions. The returns of all other hedge fund styles are not significantly impacted by an impulse to FX illiquidity.

[Figure 1 about here.]

The hypothesis that hedge fund activity increases illiquidity and volatility in both the FX and equity market is investigated in Figure 2. Panel (a) shows that there is no evidence that combined hedge funds returns of all strategies impact the market variables in the two markets. The picture looks slightly different for individual strategies, which are depicted in Panels (b)–(k). For instance equity volatility responds negatively to a positive shock to the returns of dedicated short bias funds, which is in line with the theory that short selling helps to improve market efficiency and leads to lower volatility. Positive impulses to a number of hedge funds styles are followed by lower illiquidity and volatility in the FX market. This indicates that the relation between hedge fund returns and FX market variables is bidirectional, further supporting the presence of significant FX exposure of hedge funds.

There is only weak evidence in favor of the second hypothesis. Positive returns of equity market neutral funds tend to be followed by larger volatility in equity markets. A possible reason are chain reactions of complex trading algorithms as experienced during the quant crash in August 2007 (Khandani and Lo, [2007, 2011]). More recently, the flash crash on May 6th, 2010 provides further evidence for the potential risk and volatility increasing properties of computerized trading strategies that are frequently employed by equity market neutral funds. Furthermore, global macro hedge fund returns tend to trigger a positive reaction in FX illiquidity, but this effect is not statistically significant. While Brown, Goetzmann, and Park (2000) find no evidence that hedge funds were able to move exchange rates during the Asian currency crisis, positive shocks to global macro hedge fund returns might have an increasing effect on FX illiquidity.

Overall, the evidence for an influence of hedge fund returns on market variables is rather limited. Consequently, the second hypothesis is not supported as for most hedge fund strategies there is even some evidence for a negative effect of hedge fund returns on illiquidity and volatility with equity market neutral hedge funds being the only statistically significant exception.

[Figure 2 about here.]

[14] For instance Mancini-Griffoli and Ranaldo (2010) document covered interest rate arbitrage opportunities during the recent financial crisis.

The final hypothesis concerns cross-market dynamics in volatility and illiquidity. Impulse responses of these market variables are shown in Figure 3, corroborating the presence of bidirectional effects across the FX and equity market. Panel (a) documents that a shock to equity volatility leads to an persistent increase in FX volatility as well as less liquidity in the equity and the FX market. The responses to a shock in equity illiquidity (Panel (b)) are not particularly strong: FX volatility tends to decrease after a positive shock to equity illiquidity.

Impulse response functions to shocks in FX market variables are depicted in Panels (c) and (d). An increase in FX volatility is followed by lower FX liquidity. Similarly, there is evidence for spillover effects to equity volatility. This cross-market effect is more significant with the different ordering of endogenous variables presented in the internet appendix. The results for FX illiquidity reflect cross-market illiquidity dynamics: A shock to FX illiquidity leads to less liquidity in the equity market. Moreover, evidence for feedback effects between volatility and illiquidity in the FX market is substantiated. All in all, the presence of linkages between volatility and illiquidity in the equity and the FX market can be confirmed, indicating that volatility and illiquidity in one market have predictive power for the other market.

[Figure 3 about here.]

4.4. Relation to Trend Following Risk Factors

The analysis of hedge fund excess returns did only provide marginal evidence in favor of hedge funds increasing volatility or illiquidity. To potentially corroborate the finding that hedge funds activity, measured by returns, does not increase volatility and illiquidity it is relevant to investigate the influence of trend following strategy risk factors. These factors are constructed to mirror trend following strategies that have non-linear option-like features and profit during extreme market events. Thus, in contrast to classical linear benchmark factors such as $MKT{-}RF$, SMB, $\Delta TERM$, and $\Delta CREDIT$, the trend following factors $PTFSBD$, $PTFSCOM$, $PTFSFX$, are constructed to represent hedge fund trading activity which might impact market variables. This is particularly true in light of the fact that lookback straddles deliver the difference between the maximum and the minimum price over a certain time interval and, thus, thrive during extreme market environments.

Figure 4 shows responses of market variables to shocks in the trend following factors. Interestingly, volatility in both the equity as well as the FX market increases significantly following a shock to the trend following risk factors. Thus, the trend following strategies mirrored by the $PTFSBD$, $PTFSCOM$, $PTFSFX$ factors appear to magnify financial market volatility. Moreover, there is a similar, but weaker effect for FX illiquidity, which is partly significant depending on the ordering of endogenous variables in the impulse response analysis. In contrast, there is no significant response of illiquidity in the equity market.

[Figure 4 about here.]

Given that volatility tends to increase after shocks to the trend following factors, the question arises which hedge fund styles have the largest loading on these factors. In line with the fact that trend following is the most popular strategy of CTAs (Fung and Hsieh, [1997b, 2001]), managed future funds have a significantly positive loading on all three trend following factors (see Table I.5 in the internet appendix for detailed regression results). The results for the remaining strategies are not as unambiguous. The loading of $PTFSBD$ is negative and significant for most strategies, i.e., hedge fund returns are lower in times of positive performance of $PTFSBD$. Overall, the index for all hedge funds indicates a negative loading on $PTFSBD$, whereas it is positively related to $PTFSCOM$ and $PTFSFX$. Therefore, seen as a whole, hedge funds returns increase with the trend following factors for the commodity and foreign exchange market. Given that these factors also lead to higher volatility in the equity and the FX market, there is evidence that hedge fund trading activity, in particular of CTAs, tends to be followed by larger volatility.

5. Summary and Concluding Remarks

The role of hedge funds in financial markets is controversial and there exist various hypotheses regarding the interrelation of hedge fund returns, volatility, and illiquidity. To analyze potential causalities among these variables, this paper investigates the joint dynamics of hedge fund returns and volatility as well as illiquidity in the equity and the foreign exchange market. The results yield valuable insights for hedge fund managers, investors as well as regulators and politicians.

The empirical evidence shows that a number of hedge fund strategies responds negatively to shocks in volatility. Thus, hedge funds do not profit from higher levels of volatility in the two markets. In particular the negative relation to FX volatility is remarkable. While an exposure to FX market variables is expected for global macro or emerging market hedge funds, the results show that even hedge funds focused on equity markets, such as e.g., long/short equity funds, have significant lower returns following a shock to FX volatility. This finding has important implications for performance evaluation and compensation of managers which are largely rewarded for generating alpha. Actually, part of hedge funds profits might simply be due to FX volatility exposure which is a valuable insight for fund of fund managers. Moreover, investors and fund of funds should take this FX exposure into account when making risk management and asset allocation decisions.

In contrast, a shock to equity illiquidity has a positive impact on most hedge fund strategies, which is consistent with hedge funds being able to earn liquidity premiums during illiquid market conditions. Therefore, hedge fund managers should mainly focus on volatility exposure in their risk management process and ensure a sufficient level of funding liquidity during crisis periods to take advantage of profitable trading opportunities when they emerge.

One of the key questions in the recent political and regulatory debate is whether hedge fund activity increases volatility and illiquidity in financial markets. This paper contributes to this discussion by providing empirical evidence regarding the responses of volatility and illiquidity to shocks in hedge fund returns. While there is no clear evidence regarding illiquidity, some hedge fund styles might impact equity and FX volatility. The main positive connection between these variables is due to the trend following factors for commodities and foreign exchange with CTAs exhibiting the largest loading on these factors. Moreover, the algorithmic trading activities of equity market neutral funds might increase volatility. On the other hand, dedicated short bias returns tend to be followed by lower volatility corroborating the potential benefits of short selling. Thus, I find no evidence consistent with a destabilizing effect of short selling which is frequently assumed by politicians. Overall, the response of market variables to hedge fund returns is limited. However note that the dynamics are likely to become stronger in case new regulation, e.g., more stringent capital requirements for banks, induces growth of the shadow banking system.

Finally, the paper documents strong linkages and lead/lag relations between equity and foreign exchange volatility and illiquidity. These bidirectional spillovers indicate integration of the FX and stock market which is important for risk management. Evidence from causality tests shows that FX illiquidity and equity volatility are the main drivers in the cross-market dynamics.

The results in this paper suggest a number of avenues for future research. First, given the significant exposure to FX volatility and illiquidity of many strategies, an extension of the Fung-Hsieh seven factor model by a foreign exchange risk factor promises to yield important insights to understand and model the cross section of hedge fund returns. Part of the alpha in frequently applied factor models might be attributable to hedge funds' exposure to FX volatility and/or illiquidity risk rather than managerial skill. Second, future research regarding the hedging implications for hedge funds is warranted. The returns of numerous hedge fund strategies respond negatively to shocks in equity and FX volatility, so hedge fund returns might be stabilized and potentially increased by more effective hedging against volatility risk. The effectiveness of such hedging strategies remains to be investigated. Third, the potentially volatility increasing effect of computerized trading of trend following CTAs and equity market neutral funds should be analyzed in more detail to derive potential regulatory measures such as, for instance, instituting a minimal time a limit order must be kept in the system or a ban on stub quotes.

Appendix A. Figures and Tables

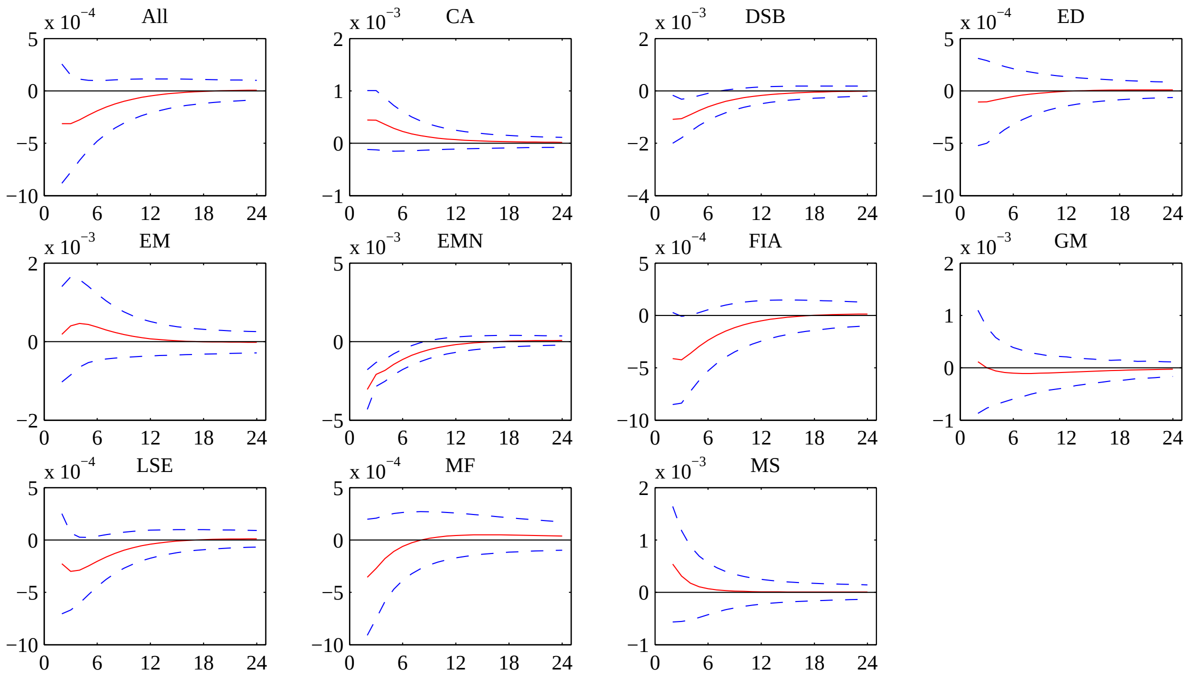

(a) Response to shock in VOL_{EQ}

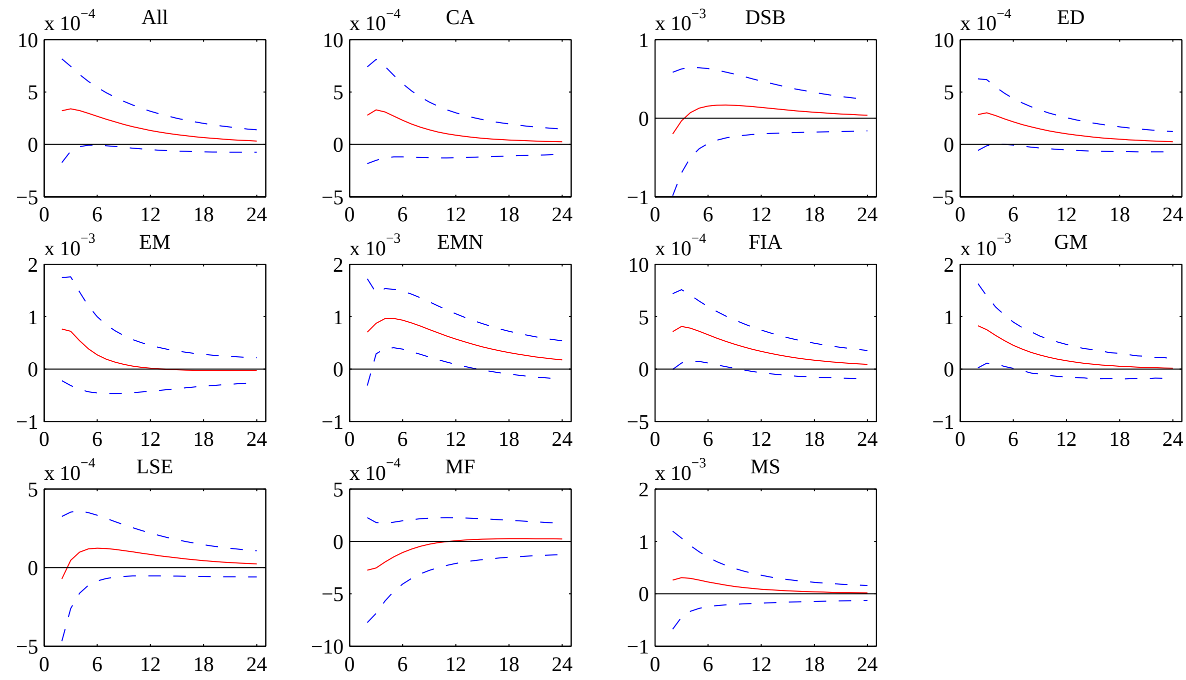

(b) Response to shock in ILL_{EQ}

Figure 1: Impulse response functions of various hedge fund returns to shocks in equity and FX volatility and illiquidity.

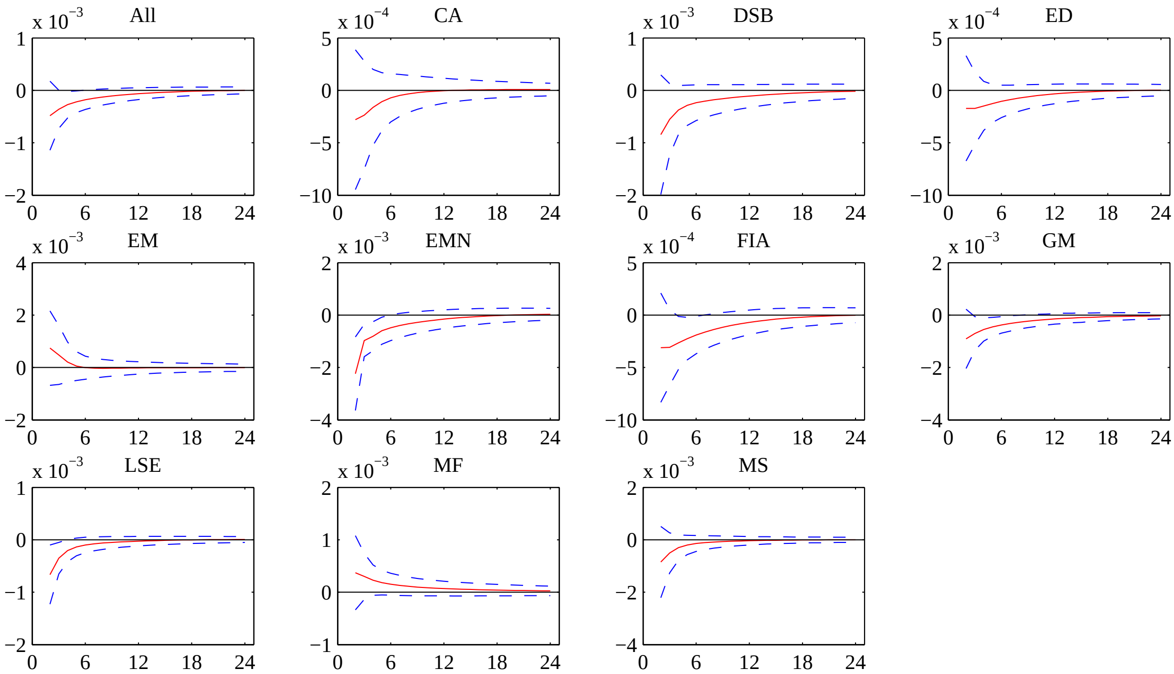

(c) Response to shock in VOL_{FX}

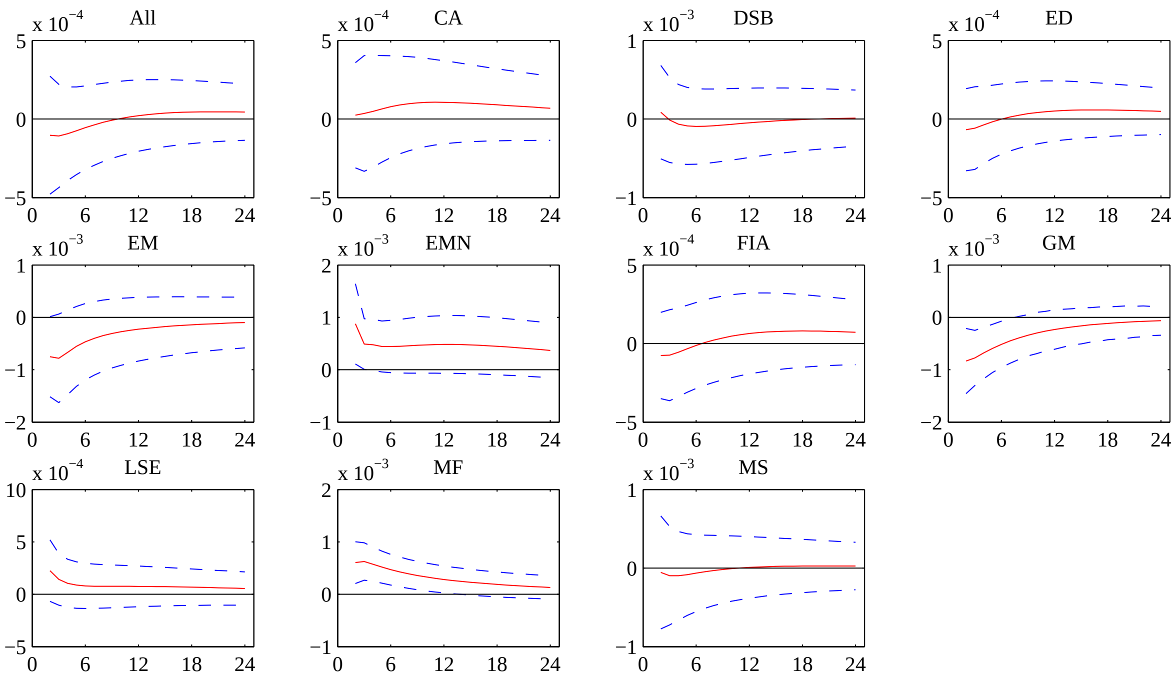

(d) Response to shock in ILL_{FX}

Figure 1 (continued): Impulse response functions of various hedge fund returns to shocks in equity and FX volatility and illiquidity.

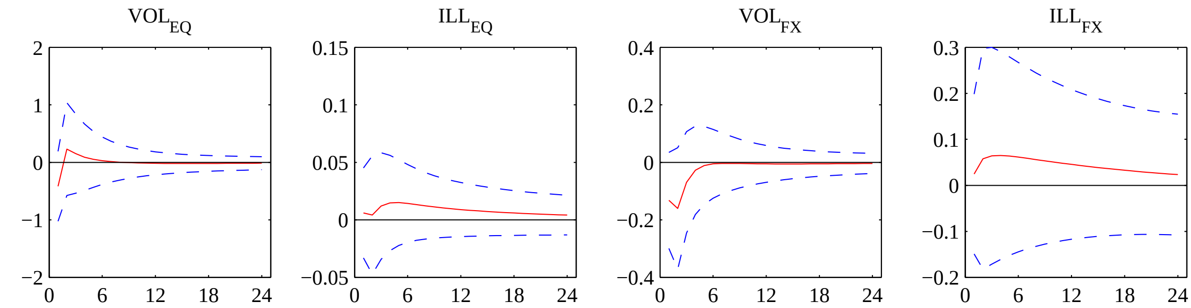

(a) Response to shock in $r_{HF,all}$

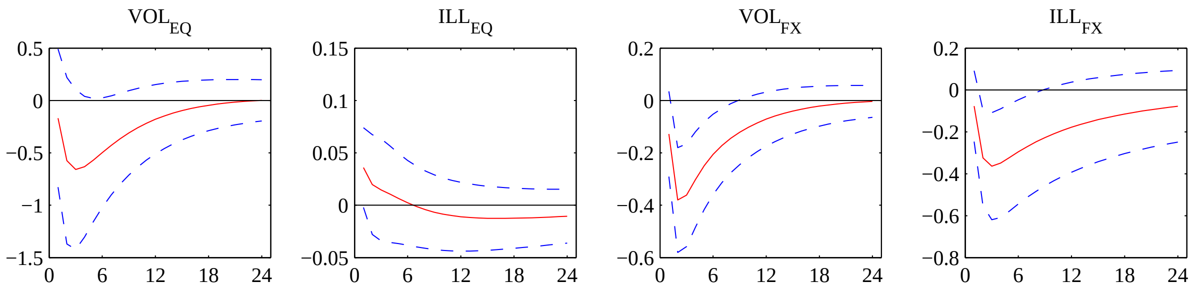

(b) Response to shock in $r_{HF,ca}$

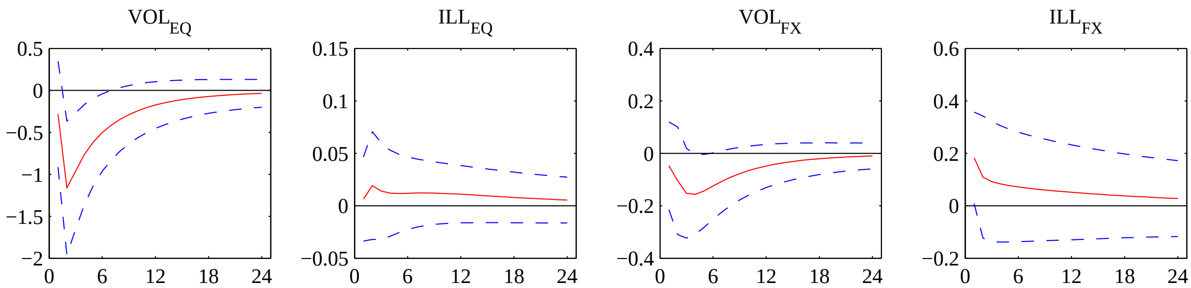

(c) Response to shock in $r_{HF,dsb}$

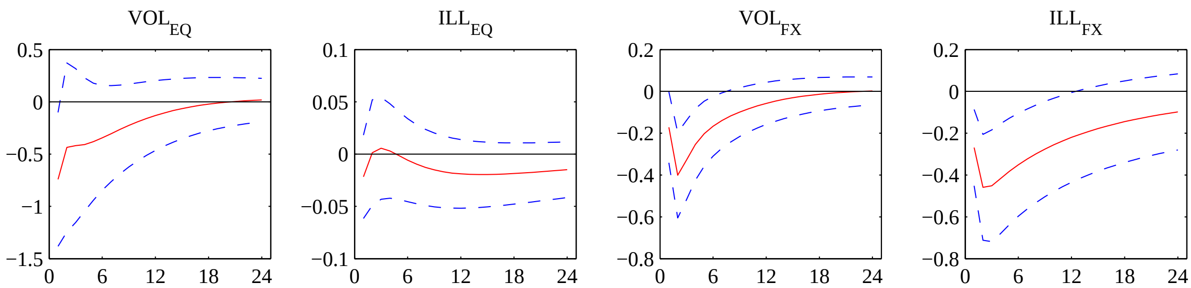

(d) Response to shock in $r_{HF,ed}$

Figure 2: Impulse response functions of equity and foreign exchange illiquidity and volatility.

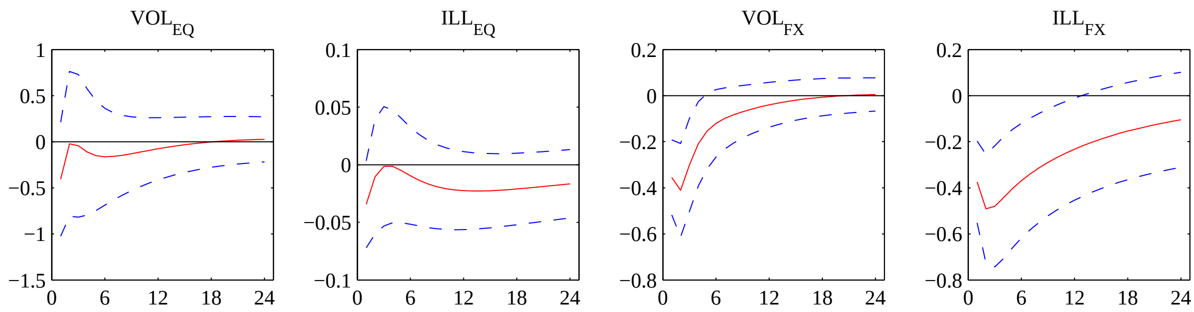

(e) Response to shock in $r_{HF,em}$

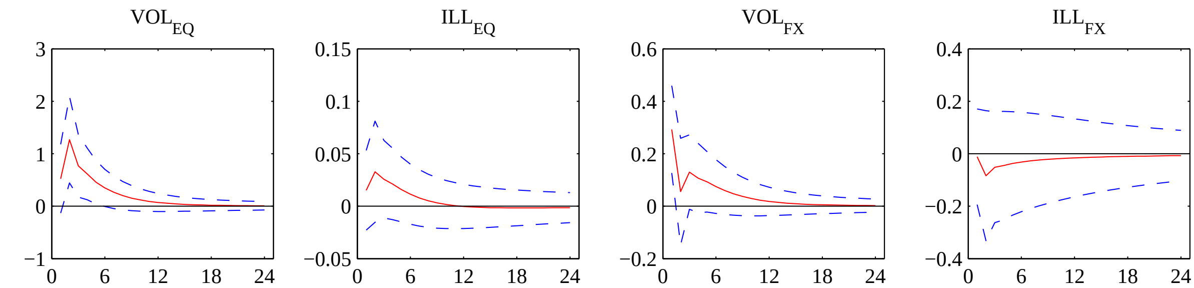

(f) Response to shock in $r_{HF,emn}$

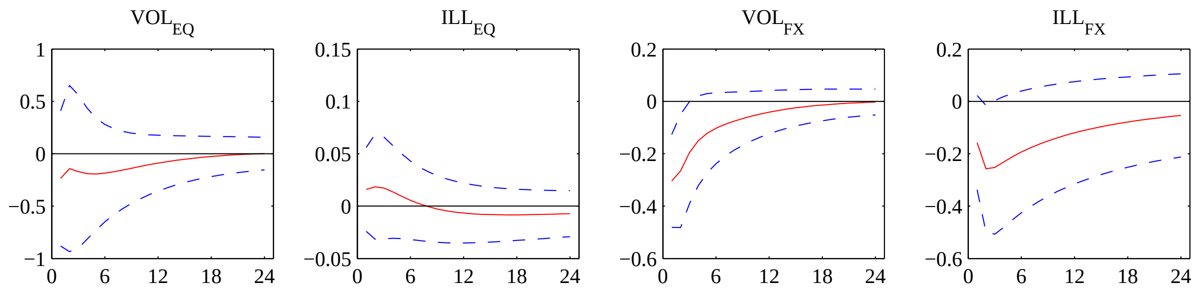

(g) Response to shock in $r_{HF,fia}$

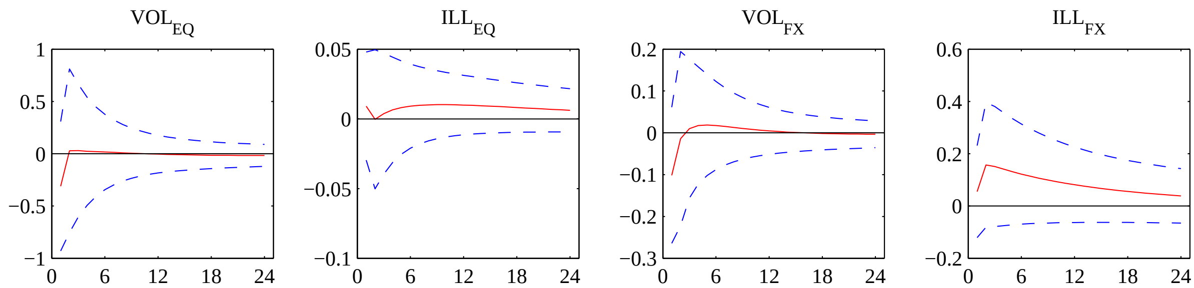

(h) Response to shock in $r_{HF,gm}$

Figure 2 (continued): Impulse response functions of equity and foreign exchange illiquidity and volatility.

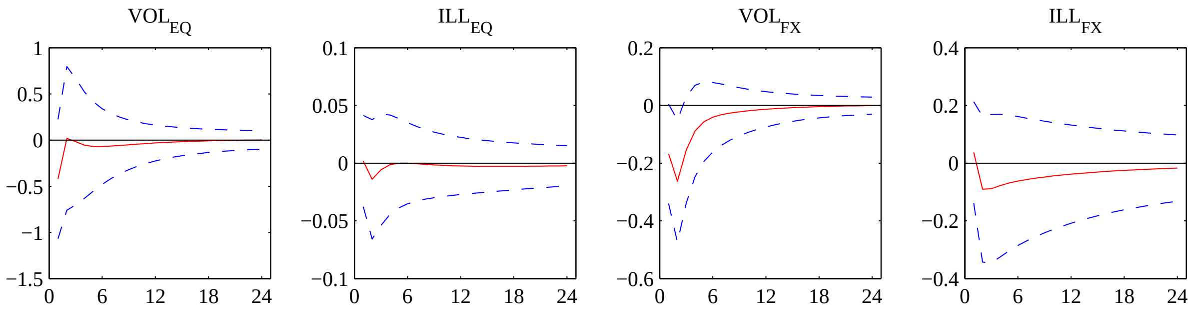

(i) Response to shock in $r_{HF,lse}$

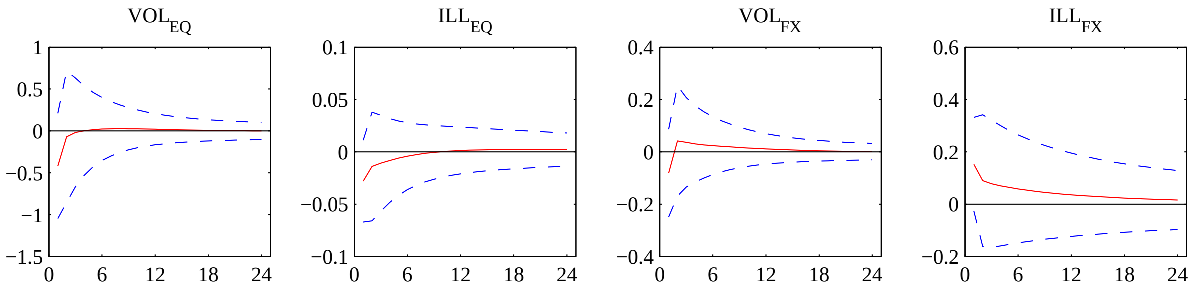

(j) Response to shock in $r_{HF,mf}$

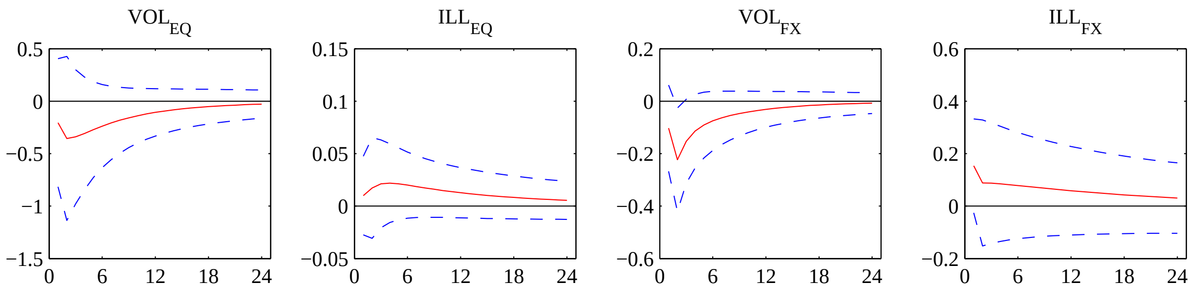

(k) Response to shock in $r_{HF,ms}$

Figure 2 (continued): Impulse response functions of equity and foreign exchange illiquidity and volatility.

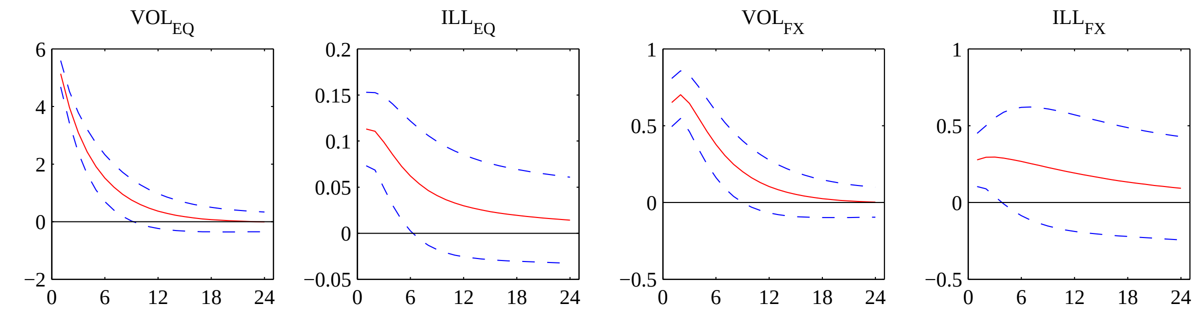

(a) Response to shock in VOL_{EQ}

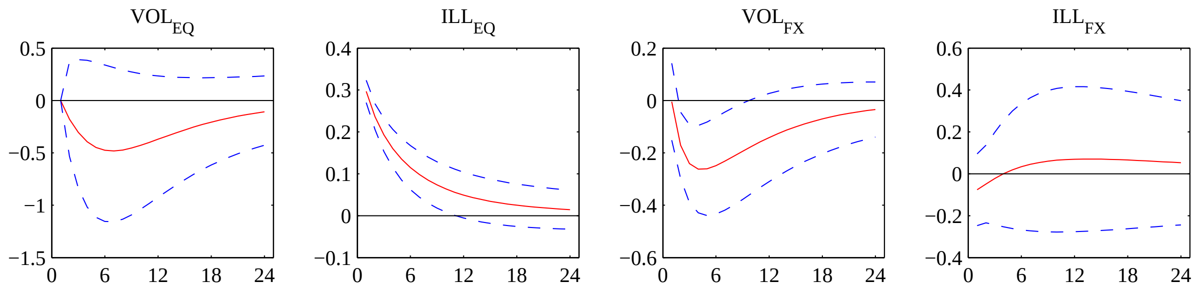

(b) Response to shock in ILL_{EQ}

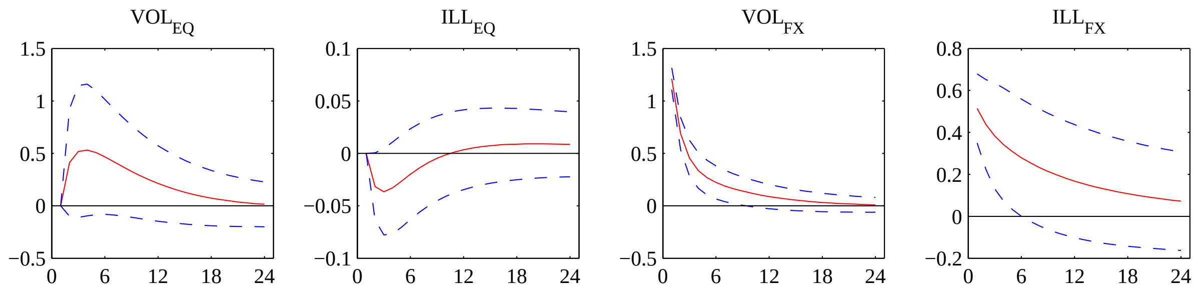

(c) Response to shock in VOL_{FX}

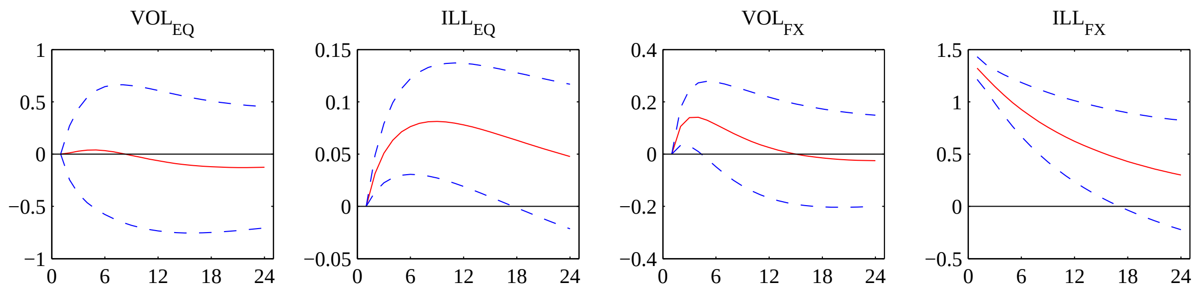

(d) Response to shock in ILL_{FX}

Figure 3: Impulse response functions of equity and foreign exchange illiquidity and volatility.

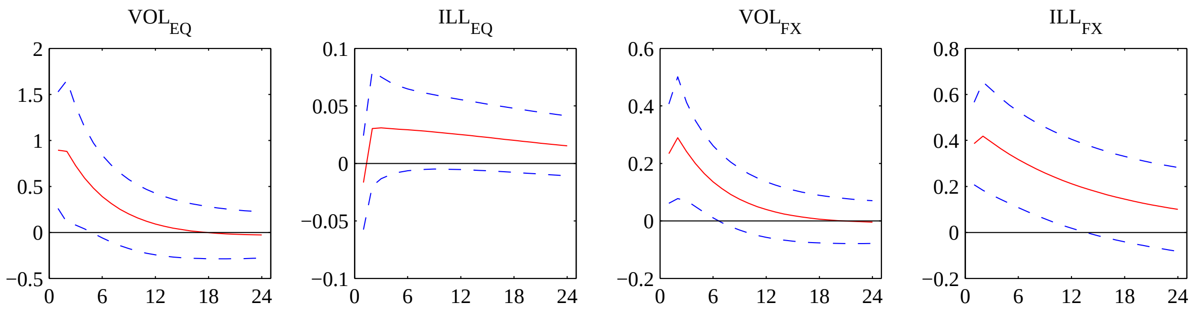

(a) Response to shock in $PTFSBD$

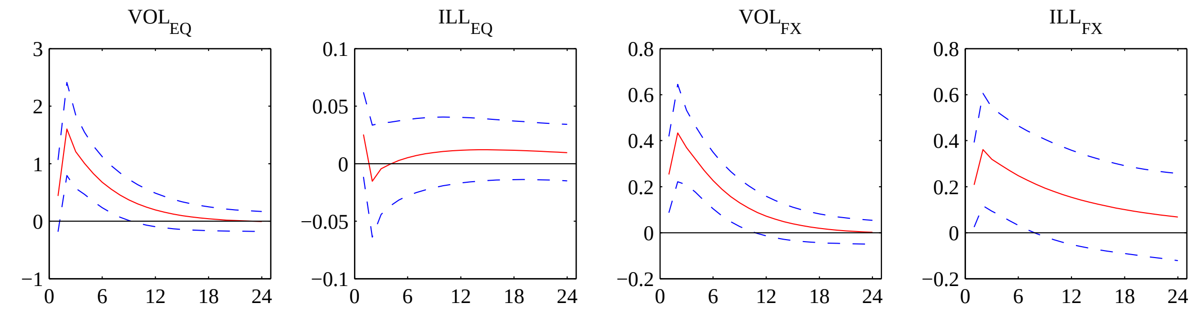

(b) Response to shock in $PTFSCOM$

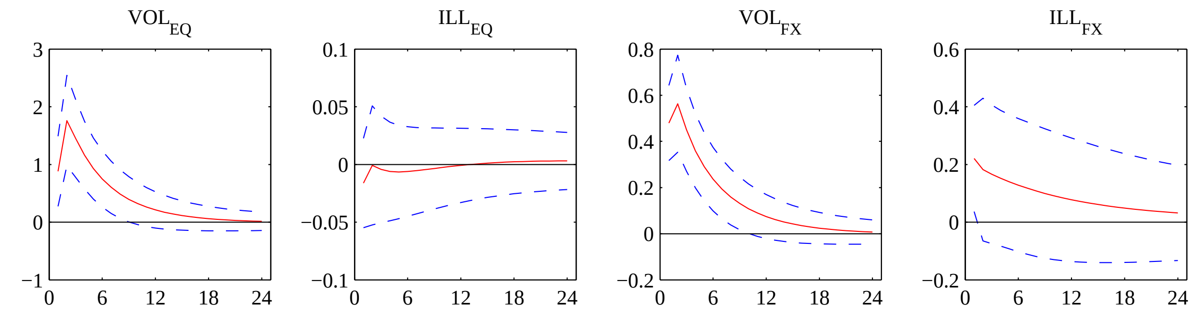

(c) Response to shock in $PTFSFX$

Figure 4: Impulse response functions of equity and foreign exchange illiquidity and volatility to trend following risk factors.

Table 1: This table shows descriptive statistics for equity and foreign exchange illiquidity and volatility measures.

	VOL_{EQ}	ILL_{EQ}	VOL_{FX}	ILL_{FX}
	(in %)	(Amihud $\times 10^5$)	(in %)	(in bps)
Mean	13.19	1.47	6.29	10.00
Median	11.30	1.52	5.95	8.88
Std. dev.	8.34	0.66	2.33	3.86
Min	3.61	0.43	2.80	5.94
Max	60.37	4.27	19.25	27.16

Table 2: This table shows pairwise correlations of excess returns of various hedge fund strategies as well as equity and foreign exchange volatility and illiquidity.

$r_{HF,all}$	1														
$r_{HF,ca}$	0.348***	1													
$r_{HF,dsb}$	0.160***	0.045	1												
$r_{HF,ed}$	0.497***	0.509***	−0.052	1											
$r_{HF,em}$	0.491***	0.194***	−0.083	0.457***	1										
$r_{HF,emn}$	0.047	−0.001	0.141**	0.042	−0.101*	1									
$r_{HF,fia}$	0.372***	0.617***	0.110*	0.294***	0.168***	0.187***	1								
$r_{HF,gm}$	0.886***	0.234***	0.208***	0.290***	0.375***	−0.078	0.319***	1							
$r_{HF,lse}$	0.676***	0.249***	−0.045	0.446***	0.357***	−0.101*	0.184***	0.407***	1						
$r_{HF,mf}$	0.352***	0.036	0.038	0.087*	0.090*	−0.007	0.081	0.343***	0.243***	1					
$r_{HF,ms}$	0.225***	0.565***	0.190***	0.272***	−0.040	0.195***	0.431***	0.117*	0.203***	0.115*	1				
VOL_{EQ}	−0.166***	0.009	−0.208***	−0.145***	−0.015	−0.258***	−0.213***	−0.085*	−0.088*	−0.034	−0.126**	1			
ILL_{EQ}	0.089*	0.216***	−0.034	0.072	−0.016	0.122**	0.149***	0.065	0.048	0.003	0.029	0.186***	1		
VOL_{FX}	−0.219***	−0.108*	−0.193***	−0.185***	−0.204***	−0.191***	−0.367***	−0.170***	−0.046	−0.060	−0.166***	0.724***	0.019	1	
ILL_{FX}	−0.041	0.004	−0.016	−0.096*	−0.229***	0.045	−0.088*	−0.143***	0.217***	0.022	0.055	0.173***	0.443***	0.370***	1

***, **, and * indicate significance at the 5%, 10%, and 25% levels, respectively.

Table 3: This table shows pair-wise Granger causality test statistics (p-values in parentheses). The null hypothesis is that the row variable does not Granger cause the column variable.

	$r_{HF,all}$	$r_{HF,ca}$	$r_{HF,dsb}$	$r_{HF,ed}$	$r_{HF,em}$	$r_{HF,emn}$	$r_{HF,fia}$	$r_{HF,gm}$	$r_{HF,lse}$	$r_{HF,mf}$	$r_{HF,ms}$
VOL_{EQ}	0.1250	1.0636	0.5005	0.1097	0.2023	4.1188***	1.0761	0.1264	0.8562	0.8668	0.1271
	(0.7237)	(0.3024)	(0.4793)	(0.7405)	(0.6528)	(0.0424)	(0.2996)	(0.7222)	(0.3548)	(0.3518)	(0.7214)
ILL_{EQ}	1.3402*	1.1686	0.2057	2.0355*	1.6492*	1.7829*	3.1198**	2.8249**	0.6445	0.2316	0.0531
	(0.2470)	(0.2797)	(0.6502)	(0.1537)	(0.1991)	(0.1818)	(0.0773)	(0.0928)	(0.4221)	(0.6304)	(0.8178)
VOL_{FX}	1.0093	0.4806	1.4563*	0.2209	1.1851	8.4185***	0.7324	0.5871	0.0963	0.9044	4.4686***
	(0.3151)	(0.4881)	(0.2275)	(0.6383)	(0.2763)	(0.0037)	(0.3921)	(0.4435)	(0.7563)	(0.3416)	(0.0345)
ILL_{FX}	0.2249	0.0133	0.0600	0.1882	2.7741**	3.7482**	0.2197	4.9015***	6.4633***	0.0171	1.6722*
	(0.6354)	(0.9083)	(0.8065)	(0.6644)	(0.0958)	(0.0529)	(0.6393)	(0.0268)	(0.0110)	(0.8959)	(0.1960)
ALL	6.5811*	3.6757	8.4003**	3.9242	3.4085	53.4195***	10.8052***	10.3673***	10.2582***	1.9855	8.4171**
	(0.1598)	(0.4517)	(0.0780)	(0.4164)	(0.4919)	(0.0000)	(0.0288)	(0.0347)	(0.0363)	(0.7384)	(0.0774)

***, **, and * indicate significance at the 5%, 10%, and 25% levels, respectively.

Table 4: This table shows pair-wise Granger causality test statistics (p-values in parentheses). The null hypothesis is that the row variable does not Granger cause the column variable.

	VOL_{EQ}	ILL_{EQ}	VOL_{FX}	ILL_{FX}
$r_{HF,all}$	2.4166*	0.0161	0.3131	0.1028
	(0.1201)	(0.8992)	(0.5758)	(0.7485)
$r_{HF,ca}$	1.1861	0.2292	8.3184***	6.7529***
	(0.2761)	(0.6321)	(0.0039)	(0.0094)
$r_{HF,dsb}$	6.3865***	0.1889	0.4636	0.3334
	(0.0115)	(0.6639)	(0.4960)	(0.5637)
$r_{HF,ed}$	0.1678	1.1430	6.1431***	4.0037***
	(0.6821)	(0.2850)	(0.0132)	(0.0454)
$r_{HF,em}$	1.0678	0.5592	3.5731**	2.2682*
	(0.3014)	(0.4546)	(0.0587)	(0.1321)
$r_{HF,emn}$	5.1245***	1.6088*	1.9534*	0.4691
	(0.0236)	(0.2047)	(0.1622)	(0.4934)
$r_{HF,fia}$	0.1325	0.0002	0.4198	1.4339*
	(0.7159)	(0.9890)	(0.5170)	(0.2311)
$r_{HF,gm}$	0.6078	0.2038	0.4278	0.9994
	(0.4356)	(0.6516)	(0.5131)	(0.3175)
$r_{HF,lse}$	1.0319	0.7101	2.0347*	1.4869*
	(0.3097)	(0.3994)	(0.1537)	(0.2227)
$r_{HF,mf}$	0.4280	0.0438	0.7955	0.2124
	(0.5130)	(0.8342)	(0.3724)	(0.6449)
$r_{HF,ms}$	0.2088	0.0244	2.6293*	0.3064
	(0.6477)	(0.8759)	(0.1049)	(0.5799)

***, **, and * indicate significance at the 5%, 10%, and 25% levels, respectively.

Table 5: This table shows pair-wise Granger causality test statistics (p-values in parentheses). The null hypothesis is that the row variable does not Granger cause the column variable.

	VOL_{EQ}	ILL_{EQ}	VOL_{FX}	ILL_{FX}
VOL_{EQ}	–	2.6495*	15.9991***	0.2368
		(0.1036)	(0.0001)	(0.6266)
ILL_{EQ}	0.7496	–	8.7437***	0.1455
	(0.3866)		(0.0031)	(0.7029)
VOL_{FX}	1.4818*	4.5106***	–	0.1973
	(0.2235)	(0.0337)		(0.6569)
ILL_{FX}	0.0055	9.5995***	5.7125***	–
	(0.9408)	(0.0019)	(0.0168)	
ALL	5.2158	10.1521***	19.4735***	0.7881
	(0.2659)	(0.0379)	(0.0006)	(0.9400)

***, **, and * indicate significance at the 5%, 10%, and 25% levels, respectively.

References

Agarwal, V., and N. Naik, 2004, "Risks and Portfolio Decisions Involving Hedge Funds", *The Review of Financial Studies*, 17(1), 63–98.

Amihud, Y., 2002, "Illiquidity and Stock Returns: Cross-Section and Time-Series Effects", *Journal of Financial Markets*, 5(1), 31–56.

Anson, M., and H. Ho, 2003, "Short Volatility Strategies: Identification, Measurement, and Risk Management", *Journal of Investment Management*, 1(2), 30–43.

Aragon, G., and P. Strahan, 2009, "Hedge Funds as Liquidity Providers: Evidence from the Lehman Bankruptcy", NBER Working paper 15336, Arizona State University and Boston College.

Bank for International Settlements, 2010, "Foreign Exchange and Derivatives Market Activity in April 2010", Triennial Central Bank Survey.

Bollen, N., and R. Whaley, 2009, "Hedge Fund Risk Dynamics: Implications for Performance Appraisal", *The Journal of Finance*, 64(2), 985–1035.

Boyson, N., C. Stahel, and R. Stulz, 2010, "Hedge Fund Contagion and Liquidity Shocks", *Journal of Finance*, 65(5), 1789–1816.

Brown, S., W. Goetzmann, and J. Park, 2000, "Hedge Funds and the Asian Currency Crisis", *The Journal of Portfolio Management*, 26(4), 95–101.

Brunnermeier, M. K., and L. H. Pedersen, 2009, "Market Liquidity and Fund Liquidity", *The Review of Financial Studies*, 22(6), 2201–2238.

Campbell, J., K. Serfaty-de Medeiros, and L. Viceira, 2010, "Global Currency Hedging", *The Journal of Finance*, 65(1), 87–121.

Cao, C., Y. Chen, B. Liang, and A. W. Lo, 2010, "Can Hedge Funds Time Market Liquidity", Working paper, MIT Sloan School of Management.

Chordia, T., A. Sarkar, and A. Subrahmanyam, 2005, "An Empirical Analysis of Stock and Bond Market Liquidity", *The Review of Financial Studies*, 18(1), 85–129.

Fama, E. F., and K. R. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds", *Journal of Financial Economics*, 33(1), 3–56.

Fleming, J., C. Kirby, and B. Ostdiek, 1998, "Information and Volatility Linkages in the Stock, Bond, and Money Markets", *Journal of Financial Economics*, 49(1), 111–137.

Fung, W., and D. A. Hsieh, 1997a, "Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds", *The Review of Financial Studies*, 10(2), 275–302.

——— , 1997b, "Survivorship Bias and Investment Style in the Returns of CTAs", *Journal of Portfolio Management*, 23(1), 30–41.

——— , 2001, "The Risk in Hedge Fund Strategies: Theory and Evidence From Trend Followers", *The Review of Financial Studies*, 14(2), 313–341.

——— , 2004, "Hedge Fund Benchmarks: A Risk-based Approach", *Financial Analysts Journal*, 60(5), 65–80.

Goyenko, R. Y., C. Holden, and C. Trzcinka, 2009, "Do Liquidity Measures Measure Liquidity?", *Journal of Financial Economics*, 92(2), 153–181.

Goyenko, R. Y., and A. D. Ukhov, 2009, "Stock and Bond Market Liquidity: A Long-Run Empirical Analysis", *Journal of Financial and Quantitative Analysis*, 44(1), 189–212.

Granger, C. W. J., 1969, "Investigating Causal Relations by Econometric Models and Cross-spectral Methods", *Econometrica*, 37, 424–438.

Harris, J., 2010, "Hedge Fund Trading Volume Rises as Algorithms Become More Popular", Risk magazin, published online on July 19, 2010.

Hasanhodzic, J., and A. W. Lo, 2007, "Can Hedge-Fund Returns Be Replicated?: The Linear Case", *Journal of Investment Management*, 5(2), 5–45.

Hendershott, T., C. M. Jones, and A. J. Menkveld, 2010, "Does Algorithmic Trading Improve Liquidity?", *The Journal of Finance*, 66(1), 1–33.

International Financial Services, 2010, *IFSL Research London: Hedge Funds 2010*. London, UK.

Jaeger, L. A., and C. Wagner, 2005, "Factor Modelling and Benchmarking of Hedge Funds: Can Passive Investments in Hedge Fund Strategies Deliver?", *The Journal of Alternative Investments*, 8(3), 9–36.

Khandani, A. E., and A. W. Lo, 2007, "What Happened to the Quants in August 2007?", *Journal of Investment Management*, 5(4), 5–54.

———, 2011, "What Happened to the Quants in August 2007? Evidence from Factors and Transactions Data", *Journal of Financial Markets*, 14(1), 1–46.

Kyle, A., 1985, "Continuous Auctions and Insider Trading", *Econometrica*, 53, 1315–1335.

Lhabitant, F.-S., 2006, *Handbook of Hedge Funds*. John Wiley & Sons, Chichester, UK.

Mancini, L., A. Ranaldo, and J. Wrampelmeyer, 2010, "Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums", Swiss National Bank Working Paper 2010-03, Swiss Finance Institute and Swiss National Bank.

Mancini-Griffoli, T., and A. Ranaldo, 2010, "Limits to Arbitrage During the Crisis: Funding Liquidity constraints and Covered Interest Parity", Working paper, Swiss National Bank.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf, 2011, "Carry Trades and Global Foreign Exchange Volatility", *The Journal of Finance*, forthcoming.

Sadka, R., 2010, "Liquidity Risk and the Cross-Section of Hedge Fund Returns", *Journal of Financial Economics*, 98(1), 54–71.

Sullivan, R. N. ., 2008, "Taming Global Village Risk", *The Journal of Portfolio Management*, 34(4), 58–67.

Turner, A., 2009, "Mansion House Speech", delivered in London, UK on September 22, 2009.

World Federation of Exchanges, 2009, *2009 Annual Report and Statistics*. Paris, France.

Part III

Appendix

Supplemental Appendix A: Taking Ambiguity to Reality: Robust Agents Cannot Trust the Data Too Much

Solution to the Simple Robust Portfolio Choice Problem

The HJB corresponding to the optimization problem in (3) is given by

$$\begin{aligned} 0 \quad = \quad & \sup_{\{C_t,\pi_t\}} \inf_{\{a\}} \left\{ U(C_t) + \frac{\partial V}{\partial t} + \frac{\partial V}{\partial W} \left[W_t(r + \pi_t(\mu - r)) - C_t + \pi_t^2 \sigma^2 W_t^2 a \right] \right. \\ & \left. + \frac{1}{2} \frac{\partial^2 V}{\partial W^2} \pi_t^2 \sigma^2 W_t^2 + \frac{1}{2\Psi} (\pi_t \sigma W_t a)^2 \right\}. \qquad \text{(I.1)} \end{aligned}$$

Minimizing (I.1) with respect to a and replacing $a^* = -\frac{\partial V}{\partial W}\Psi$ in the HJB gives[15]

$$\begin{aligned} 0 \quad = \quad & \sup_{\{C_t,\pi_t\}} \left\{ \frac{C_t^{1-\gamma}}{1-\gamma} - \beta V + \frac{\partial V}{\partial t} + \frac{\partial V}{\partial W} \left[W_t(r + \pi_t(\mu - r)) - C_t \right] + \frac{1}{2} \frac{\partial^2 V}{\partial W^2} \pi_t^2 \sigma^2 W_t^2 \right. \\ & \left. - \frac{1}{2} \Psi \left(\frac{\partial V}{\partial W} \pi_t \sigma W_t \right)^2 \right\}, \qquad \text{(I.2)} \end{aligned}$$

implying the following two first order conditions for optimal consumption and portfolio choice

$$C_t^* \quad = \quad \left(\frac{\partial V}{\partial W} \right)^{-\frac{1}{\gamma}} \qquad \text{(I.3)}$$

$$\pi_t^* \quad = \quad - \frac{\frac{\partial V}{\partial W}}{\left(\frac{\partial^2 V}{\partial W^2} - \Psi \left(\frac{\partial V}{\partial W} \right)^2 \right) W_t} \frac{\mu - r}{\sigma^2}, \qquad \text{(I.4)}$$

where (I.3) is independent of the concern for robustness and (I.4) collapses to the standard solution for $\Psi \to 0$. Choosing

$$\Psi(W,t) = \frac{\vartheta}{(1-\gamma)V(W,t)}, \qquad \text{(I.5)}$$

where ϑ measures the degree of robustness, CRRA preferences become homothetic and the value function can be shown to be of the form $V(W,t) = g(t,T)^{\gamma} \frac{W^{1-\gamma}}{1-\gamma}$. Substituting this guess and Ψ along with the optimal values C_t^* and π_t^* into (I.2), yields an ordinary differential equation for $g(t,T)$

$$0 = 1 - g(t,T) \frac{1}{\gamma} \left(\beta - r(1-\gamma) - \frac{1-\gamma}{2(\gamma + \vartheta)} \left(\frac{\mu - r}{\sigma} \right)^2 \right) + \frac{\partial g}{\partial t}, \qquad \text{(I.6)}$$

[15]Explicitly taking the time preference into account, introduces the additional βV term into the HJB, where β measures the agent's time preference.

whose solution is given by

$$g(t,T)=\frac{1}{\chi}+e^{\chi\cdot(T-t)}\cdot\text{constant}, \tag{I.7}$$

where $\chi=\frac{1}{\gamma}\left(\beta-r(1-\gamma)-\frac{1-\gamma}{2(\gamma+\vartheta)}\left(\frac{\mu-r}{\sigma}\right)^2\right)$. From the boundary condition $V(W,T)=0$ it follows that $g(T,T)=0$ and, hence, the constant term in (I.7) must be equal to $-1/\chi$, giving the following value function

$$V(W,t)=\left(\frac{1-e^{\chi\cdot(T-t)}}{\chi}\right)^{\gamma}\frac{W^{1-\gamma}}{1-\gamma} \tag{I.8}$$

and associated optimal consumption and portfolio choices as given in Equation (4).

Robustness in an Economy with Rare Events

Solution to the Robust Portfolio Choice Problem

The HJB corresponding to the optimization problem in (21) is given by

$$\begin{aligned}0 \;=\; \sup_{\{C_t,\pi_t\}}\inf_{\{a,b\}} \Bigg\{ & \frac{C_t^{1-\gamma}}{1-\gamma}-\beta V+\frac{\partial V}{\partial t}+\frac{\partial V}{\partial W}\left[W_t(r+\pi_t(\mu_S-r))-C_t\right]+\frac{1}{2}\frac{\partial^2 V}{\partial W^2}\pi_t^2\sigma^2W_t^2 \\ & +\lambda_Y e^a V\left(E^{\{a,b\}}\left[(1+(e^{\xi^Y}-1)\pi_t)^{1-\gamma}\right]-1\right)+\frac{1}{\Psi}H(a,b)\Bigg\}.\end{aligned} \tag{I.9}$$

Guessing the value function to be $V(W,t)=g(t,T)^{\gamma}\frac{W_t^{1-\gamma}}{1-\gamma}$, the first order conditions of the agent's problem are

$$C_t^* = \frac{W_t}{g(t,T)} \tag{I.10}$$

$$0 = (\mu_S-r)-\gamma\sigma^2\pi_t^*+\lambda_Y e^a E^{\{b\}}\left[\left(1+(e^{\xi^Y}-1)\pi_t^*\right)^{-\gamma}\left(e^{\xi^Y}-1\right)\right] \tag{I.11}$$

$$0 = E^{\{b\}}\left[\left(1+(e^{\xi^Y}-1)\pi_t\right)^{1-\gamma}\right]-1+\frac{1-\gamma}{\vartheta}\left(a^*+\frac{1}{2}b^2\sigma_Y^2+2d\left(e^{a^*+b^2\sigma_Y^2}-1\right)\right) \tag{I.12}$$

$$0 = \frac{\partial\left(E^{\{b^*\}}\left[\left(1+(e^{\xi^Y}-1)\pi_t\right)^{1-\gamma}\right]\right)}{\partial b}+\frac{1-\gamma}{\vartheta}\left(b^*\sigma_Y^2\left(1+2de^{a+(b^*)^2\sigma_Y^2}\right)\right) \tag{I.13}$$

for C_t^*, π_t^*, a^*, and b^*, respectively. In order to compute the expectations and the derivative thereof, recall that $\xi^Y\sim N(\mu_Y,\sigma_Y^2)$ together with $E^{\{b\}}\left[e^{\xi^Y}\right]=e^{\mu_Y+\frac{1}{2}\sigma_Y^2}\cdot e^{b\sigma_Y^2}$ implies $\xi^Y_{\{b\}}\sim N(\mu_Y,\sigma_Y^2(1+2b))$.[16] From here, the optimal decision rule together with the guess, substituted into (I.9), implies an ordinary differential equation for $g(t,T)$

$$\begin{aligned}0 \;=\; & g(t,T)\Bigg[(1-\gamma)\left(r+\pi_t^*(\mu_s-r)-\frac{1}{2}\gamma(\pi_t^*\sigma)^2\right)-\beta+\lambda_Y e^{a^*}\left(E^{\{b^*\}}\left[(1+(e^{\xi^Y}-1)\pi_t^*)^{1-\gamma}\right]-1\right) \\ & +\frac{\lambda_Y(1-\gamma)}{\vartheta}\left(1+d+e^{a^*}\left(a^*+\frac{1}{2}(b^*\sigma_Y)^2-1+d\left(e^{a^*+(b^*\sigma_Y)^2}-2\right)\right)\right)\Bigg] \\ & +\frac{\partial g(t,T)}{\partial t}\gamma+\gamma,\end{aligned} \tag{I.14}$$

[16]The derivative follows from the definition of the expectation:

$$\frac{\partial\left(E^{\{b^*\}}\left[\left(1+(e^{\xi^Y}-1)\pi_t\right)^{1-\gamma}\right]\right)}{\partial b}=E^{\{b^*\}}\left[\left(1+(e^{\xi^Y}-1)\pi_t\right)^{1-\gamma}\left(\left(\frac{\xi^Y-\mu_Y}{\sigma_Y(1+2b)}\right)^2-(1+2b)^{-1}\right)\right].$$

whose solution is given by $g(t,T) = -\frac{\gamma}{\chi} + e^{-\frac{\chi}{\gamma}\cdot(T-t)} \cdot \left(1 + \frac{\gamma}{\chi}\right)$ when $V(W,T) = \frac{W_T^{1-\gamma}}{1-\gamma}$ and χ represents the coefficient multiplying $g(t,T)$ in (I.14).

Market Equilibrium

In the equilibrium of the considered production economy, the representative agent holds the stock, such that $\pi_t^* = 1$, which implies that a^* and b^* are implicitly defined by simplifying Equations (I.12) and (I.13). The former reduces to

$$0 = E^{\{b^*\}}\left[\left(e^{\xi^Y}\right)^{1-\gamma}\right] - 1 + \frac{1-\gamma}{\vartheta}\left(a^* + \frac{1}{2}(b^*)^2\sigma_Y^2 + 2d\left(e^{a^*+(b^*)^2\sigma_Y^2} - 1\right)\right). \tag{I.15}$$

From the distribution of ξ^Y under the alternative, it follows that $\xi^Y_{\{b\}}(1-\gamma) \sim N(\mu_Y(1-\gamma), \sigma_Y^2(1+2b)(1-\gamma)^2)$, so that Equation (I.15) becomes

$$0 = e^{\mu_Y(1-\gamma)+\frac{1}{2}\sigma_Y^2(1+2b^*)(1-\gamma)^2} - 1 + \frac{1-\gamma}{\vartheta}\left(a^* + \frac{1}{2}(b^*)^2\sigma_Y^2 + 2d\left(e^{a^*+(b^*)^2\sigma_Y^2} - 1\right)\right). \tag{I.16}$$

Similarly, Equation (I.13) is then given by

$$0 = e^{\mu_Y(1-\gamma)+\frac{1}{2}\sigma_Y^2(1+2b^*)(1-\gamma)^2}\sigma_Y^2(1-\gamma)^2 + \frac{1-\gamma}{\vartheta}\left(b^*\sigma_Y^2\left(1 + 2de^{a^*+(b^*)^2\sigma_Y^2}\right)\right). \tag{I.17}$$

The equilibrium values a^*, b^*, and $\pi_t = 1$, together with Equation (I.11), imply the following condition for the equity premium of the economy

$$\begin{aligned} \mu_S - r &= \gamma\sigma^2 - \lambda_Y e^{a^*} E^{\{b^*\}}\left[e^{\xi^Y(1-\gamma)} - e^{-\gamma\xi^Y}\right] \\ \Leftrightarrow \mu_S - r &= \gamma\sigma^2 - \lambda_Y e^{a^*}\left(e^{\mu_Y(1-\gamma)+\frac{1}{2}(1+\gamma^2)\sigma_Y^2(1+2b^*)} - e^{-\gamma\mu_Y+\frac{1}{2}\sigma_Y^2\gamma^2(1+2b^*)}\right). \end{aligned} \tag{I.18}$$

The equity premium, together with the stochastic discount factor (SDF) of the economy can then be used to derive the drift of the asset price process for pricing purposes. Based on the pricing kernel m_t of the economy

$$dm_t = \left(-r - \lambda e^a E^{\{b\}}\left[e^{-\gamma\xi^Y} - 1\right]\right) m_t dt - \gamma\sigma m_t dB_t + m_t\left[e^{-\gamma\xi^Y} - 1\right] dN_t, \tag{I.19}$$

the continuous time pricing equation $E\left[d\left(m_t C_t\right)\right] = 0$, where C_t is the call price, gives

$$rC_t dt = \frac{\partial C}{\partial t}dt + \left(\mu_S - \gamma\sigma^2\right) S_t \frac{\partial C}{\partial S}dt + \frac{1}{2}\frac{\partial^2 C}{\partial S^2}\sigma^2 S_t^2 dt + \lambda_Y e^a E^{\{b\}}\left[e^{-\gamma\xi^Y}\left(C_t - C_{t-}\right)\right] dt. \tag{I.20}$$

Using the equity premium (Equation (I.18)) gives the partial differential equation that the call has to satisfy[17]

$$rC_t = \frac{\partial C}{\partial t} + \left(r - \lambda_Y e^a E^{\{b\}}\left[e^{-\gamma\xi^Y}\left(e^{\xi^Y} - 1\right)\right]\right) S_t \frac{\partial C}{\partial S} + \frac{1}{2}\frac{\partial^2 C}{\partial S^2}\sigma^2 S_t^2 + \lambda_Y e^a E^{\{b\}}\left[e^{-\gamma\xi^Y}\left(C_t - C_{t-}\right)\right]. \tag{I.21}$$

From here the risk neutral drift, required for pricing the call, is given by the coefficient multiplying the $\partial C/\partial S$ term.

[17]Note, there is no distinction between risk neutral and physical intensity and jump size being made here.

Solution to the Portfolio Problem in an Economy with Return Predictability

Denoting by $\bar{x}$, and $\bar{\mu}$ the means of the predictor variable and instantaneous drift in (37), respectively, the asset dynamics can be rewritten according to

$$dS_t = (\bar{\mu} + \nu(X(t) - \bar{x}))\, S_t dt + \sigma S_t dB_t^{(1)}. \tag{I.22}$$

An agent with ambiguity aversion who wishes to maximize terminal wealth, faces the following portfolio optimization problem

$$\sup_{\{\pi_t\}} \inf_{\{a\}} E^a \left[U(W_T) + \int_0^T \frac{1}{2\Psi} (\pi_t \sigma W_t a)^2 dt \right] \tag{I.23}$$

subject to

$$dW_t = \left[W_t \left(r + \pi_t \left([\bar{\mu} + \nu(X(t) - \bar{x})] - r \right) \right) + \pi_t^2 \sigma^2 W_t^2 a \right] dt + \pi_t \sigma W_t dB_t^{(1)},$$

in accordance with the setup in Equation (2). The corresponding HJB is given by

$$\begin{aligned} 0 \;=\; \sup_{\{\pi_t\}} \inf_{\{a\}} \Bigg\{ & -\beta V + \frac{\partial V}{\partial t} + \frac{\partial V}{\partial W} \left[W_t (r + \pi_t(\mu(X,t) - r)) + \pi_t^2 \sigma^2 W_t^2 a \right] \\ & + \frac{1}{2} \frac{\partial^2 V}{\partial W^2} \pi_t^2 \sigma^2 W_t^2 + \frac{1}{2\Psi} (\pi_t \sigma W_t a)^2 + \frac{1}{2} \frac{\partial^2 V}{\partial X^2} b^2 \\ & + \frac{\partial^2 V}{\partial W \partial X} \pi_t \sigma W_t b \rho + \frac{\partial V}{\partial X} \kappa_X (\bar{x} - X(t)) \Bigg\}. \end{aligned} \tag{I.24}$$

Minimizing (I.24) with respect to a and replacing $a^* = -\frac{\partial V}{\partial W} \Psi$ in the HJB gives

$$\begin{aligned} 0 \;=\; \sup_{\{\pi_t\}} \Bigg\{ & -\beta V + \frac{\partial V}{\partial t} + \frac{\partial V}{\partial W} \left[W_t (r + \pi_t(\mu(X,t) - r)) \right] \\ & + \frac{1}{2} \frac{\partial^2 V}{\partial W^2} \pi_t^2 \sigma^2 W_t^2 - \frac{1}{2} \Psi \left(\frac{\partial V}{\partial W} \pi_t \sigma W_t \right)^2 + \frac{\partial V}{\partial X} \kappa_X (\bar{x} - X_t) \\ & + \frac{1}{2} \frac{\partial^2 V}{\partial X^2} b^2 + \frac{\partial^2 V}{\partial W \partial X} \pi_t \sigma W_t b \rho \Bigg\}, \end{aligned} \tag{I.25}$$

implying the following first order condition for optimal portfolio choice

$$\pi_t^* \;=\; -\frac{\frac{\partial V}{\partial W}}{\left(\frac{\partial^2 V}{\partial W^2} - \Psi \left(\frac{\partial V}{\partial W} \right)^2 \right) W_t} \frac{\mu(X,t) - r}{\sigma^2} - \frac{\frac{\partial^2 V}{\partial W \partial X}}{\left(\frac{\partial^2 V}{\partial W^2} - \Psi \left(\frac{\partial V}{\partial W} \right)^2 \right) W_t} \frac{b \rho}{\sigma}. \tag{I.26}$$

For utility functions, such as $U(W) = \frac{W^{1-\gamma}}{1-\gamma}$, where the value function is separable in wealth and for $\Psi(W, X, t) = \frac{\vartheta}{(1-\gamma) V(W,X,t)}$ the value function satisfies

$$V(W, X, t) = g(X, t) \cdot U(W). \tag{I.27}$$

Thus, (I.26) can be written more compactly as

$$\pi_t^* \;=\; \frac{1}{\gamma + \vartheta} \left(\frac{\mu(X,t) - r}{\sigma^2} + \frac{\frac{\partial g}{\partial X}}{g(X,t)} \frac{\sigma b \rho}{\sigma^2} \right). \tag{I.28}$$

Substituting the partial derivatives of $V(W,X,t)$ into (I.25) gives a partial differential equation for $g(X,t)$

$$
\begin{aligned}
0 &= \frac{\partial g}{\partial t} + g(X,t)\left((1-\gamma)(r+\pi_t^*(\mu(X,t)-r)) - \beta - \frac{1}{2}(1-\gamma)(\pi_t^*)^2\sigma^2(\gamma+\vartheta)\right) \\
&\quad + \frac{\partial g}{\partial X}\left(\kappa_X(\bar{x}-X(t)) + (1-\gamma)\pi_t^*\sigma b\rho\right) + \frac{1}{2}\frac{\partial^2 g}{\partial X^2}b^2, \qquad \text{(I.29)}
\end{aligned}
$$

which simplifies to

$$
0 = \frac{\partial g}{\partial t} + \frac{\partial g}{\partial X}\kappa_X(\bar{x}-X(t)) + \frac{1}{2}\frac{\partial^2 g}{\partial X^2}b^2 + \left((1-\gamma)r - \beta + \frac{1}{2}\sigma^2(1-\gamma)(\gamma+\vartheta)(\pi_t^*)^2\right)g(X,t). \qquad \text{(I.30)}
$$

Using the guess, in accordance with Kim and Omberg (1996),

$$
g(X,t) = \exp\left\{A(\tau) + B(\tau)[\mu(X,t)-r] + \frac{1}{2}C(\tau)[\mu(X,t)-r]^2\right\} \qquad \text{(I.31)}
$$

with $\tau = T-t$ in (I.30) yields three ordinary differential equations (ODEs) for the coefficient functions C, B, and A.

$$
\begin{aligned}
\frac{\partial C}{\partial t} &= c_1 C(\tau)^2 + c_2 C(\tau) + c_3 && \text{(I.32)} \\
\frac{\partial B}{\partial t} &= c_1 C(\tau)B(\tau) + \frac{1}{2}c_2 B(\tau) + \kappa_X(\bar{\mu}-r)C(\tau) && \text{(I.33)} \\
\frac{\partial A}{\partial t} &= \kappa_X(\bar{\mu}-r)B(\tau) + \frac{1}{2}c_1 B(\tau)^2 + \frac{1}{2}\nu^2 b^2 C(\tau) + (1-\gamma)r - \beta && \text{(I.34)}
\end{aligned}
$$

where

$$
\begin{aligned}
c_1 &= \nu^2 b^2\left(1 + \frac{1-\gamma}{\gamma+\vartheta}\rho^2\right) && \text{(I.35)} \\
c_2 &= 2\left(\frac{1-\gamma}{\gamma+\vartheta}\frac{b\rho\nu}{\sigma} - \kappa_X\right) && \text{(I.36)} \\
c_3 &= \frac{1-\gamma}{\gamma+\vartheta}\frac{1}{\sigma^2}. && \text{(I.37)}
\end{aligned}
$$

For reasonable values of ν, γ and ϑ, Kim and Omberg (1996) show that the ODEs admit well-behaved solutions of the following form

$$
\begin{aligned}
C(\tau) &= \frac{2c_3(1-e^{-\sqrt{q}\tau})}{(\sqrt{q}-c_2) + (\sqrt{q}+c_2)e^{-\sqrt{q}\tau}} && \text{(I.38)} \\
B(\tau) &= \frac{4\sqrt{q}c_3\kappa_X(\bar{\mu}-r)(1-e^{-\sqrt{q}\tau/2})^2}{q\left((\sqrt{q}-c_2) + (\sqrt{q}+c_2)e^{-\sqrt{q}\tau}\right)} && \text{(I.39)} \\
A(\tau) &= \left(c_3\left(\frac{2\kappa_X^2(\bar{\mu}-r)^2}{q} + \frac{\nu^2 b^2}{\sqrt{q}-c_2}\right) + r(1-\gamma) - \beta\right)\tau \\
&\quad + \frac{4c_3\kappa_X^2(\bar{\mu}-r)^2\left((2c_2+\sqrt{q})e^{-\sqrt{q}\tau} - 4c_2 e^{-\sqrt{q}\tau/2} + 2c_2 - \sqrt{q}\right)}{q\sqrt{q}\left((\sqrt{q}-c_2) + (\sqrt{q}+c_2)e^{-\sqrt{q}\tau}\right)} \\
&\quad + \frac{2c_3\nu^2 b^2}{q-c_2^2}\ln\left|\frac{(\sqrt{q}-c_2) + (\sqrt{q}+c_2)e^{-\sqrt{q}\tau}}{2\sqrt{q}}\right|, && \text{(I.40)}
\end{aligned}
$$

where $q = c_1^2 - 4c_1c_3$.

Equivalence between Continuous-time and Discrete-time Model in an Economy with Return Predictability

This section proofs the equivalence between Equation (35) and the continuous-time VAR model in (38) and (I.22). First, applying Itô's Lemma, the continuous-time dynamics are transformed to log-returns

$$\begin{bmatrix} d(\log S_t + 0.5\sigma^2 t - \bar{\mu} t) \\ d(X_t - \bar{x}) \end{bmatrix} = \begin{bmatrix} 0 & \nu \\ 0 & -\kappa_X \end{bmatrix} \cdot \begin{bmatrix} \log S_t + 0.5\sigma^2 t - \bar{\mu} t \\ X_t - \bar{x} \end{bmatrix} + \begin{bmatrix} \sigma & 0 \\ b\rho & b\sqrt{1-\rho^2} \end{bmatrix} \cdot \begin{bmatrix} dB_t^1 \\ dB_t^2 \end{bmatrix}. \tag{I.41}$$

Based on the results of Bergstrom (1984), the discrete-time process implied by Equation (I.41) is:

$$\begin{aligned} &\begin{bmatrix} \log S_{t+\Delta} + 0.5\sigma^2 (t+\Delta) - \bar{\mu}(t+\Delta) \\ X_{t+\Delta} - \bar{x} \end{bmatrix} \\ &\quad = \exp\left\{ \begin{bmatrix} 0 & \nu \\ 0 & -\kappa_X \end{bmatrix} \Delta \right\} \cdot \begin{bmatrix} \log S_t + 0.5\sigma^2 t - \bar{\mu} t \\ X_t - \bar{x} \end{bmatrix} + \begin{bmatrix} u_{1,t+\Delta} \\ u_{2,t+\Delta} \end{bmatrix}, \end{aligned} \tag{I.42}$$

where

$$\begin{bmatrix} u_{1,t+\Delta} \\ u_{2,t+\Delta} \end{bmatrix} = \int_0^\Delta \exp\left\{ \begin{bmatrix} 0 & \nu \\ 0 & -\kappa_X \end{bmatrix} (\Delta - \tau) \right\} \begin{bmatrix} \sigma & 0 \\ b\rho & b\sqrt{1-\rho^2} \end{bmatrix} \begin{bmatrix} dB_{t+\tau}^1 \\ dB_{t+\tau}^2 \end{bmatrix}. \tag{I.43}$$

The solution for the exponential term in Equation (I.42) can be derived using a similar proof as in Campbell, Chacko, Rodriguez, and Viceira (2004):

$$\exp\left\{ \begin{bmatrix} 0 & \nu \\ 0 & -\kappa_X \end{bmatrix} \Delta \right\} = \begin{bmatrix} 1 & \frac{\nu}{\kappa_X}\left(1 - e^{-\kappa_X \Delta}\right) \\ 0 & e^{-\kappa_X \Delta} \end{bmatrix}. \tag{I.44}$$

Consequently,

$$\begin{aligned} &\begin{bmatrix} \log S_{t+\Delta} + 0.5\sigma^2 (t+\Delta) - \bar{\mu}(t+\Delta) \\ X_{t+\Delta} - \bar{x} \end{bmatrix} \\ &\quad = \begin{bmatrix} 1 & \frac{\nu}{\kappa_X}\left(1 - e^{-\kappa_X \Delta}\right) \\ 0 & e^{-\kappa_X \Delta} \end{bmatrix} \cdot \begin{bmatrix} \log S_t + 0.5\sigma^2 t - \bar{\mu} t \\ X_t - \bar{x} \end{bmatrix} + \begin{bmatrix} u_{1,t+\Delta} \\ u_{2,t+\Delta} \end{bmatrix}, \end{aligned} \tag{I.45}$$

where

$$\begin{bmatrix} u_{1,t+\Delta} \\ u_{2,t+\Delta} \end{bmatrix} = \int_0^\Delta \begin{bmatrix} 1 & \frac{\nu}{\kappa_X}\left(1 - e^{-\kappa_X (\Delta-\tau)}\right) \\ 0 & e^{-\kappa_X (\Delta-\tau)} \end{bmatrix} \cdot \begin{bmatrix} \sigma & 0 \\ b\rho & b\sqrt{1-\rho^2} \end{bmatrix} \begin{bmatrix} dB_{t+\tau}^1 \\ dB_{t+\tau}^2 \end{bmatrix}. \tag{I.46}$$

The variance-covariance matrix of the error terms is given by

$$\text{Var} \begin{bmatrix} u_{1,t+\Delta} \\ u_{2,t+\Delta} \end{bmatrix} = \int_0^\Delta \begin{bmatrix} D_{11} & D_{12} \\ D_{12} & D_{22} \end{bmatrix} d\tau, \tag{I.47}$$

with

$$\begin{aligned} D_{11} &= \sigma^2 + \frac{2\rho\sigma b\nu}{\kappa_X}\left(1 - e^{-\kappa_X(\Delta-\tau)}\right) + \frac{\nu^2 b^2}{\kappa_X^2}\left(1 - e^{-\kappa_X(\Delta-\tau)}\right)^2, \\ D_{12} &= \rho\sigma b e^{-\kappa_X(\Delta-\tau)} + \frac{b^2\nu}{\kappa_X}\left(e^{-\kappa_X(\Delta-\tau)} - e^{-2\kappa_X(\Delta-\tau)}\right), \\ D_{22} &= b^2 e^{-2\kappa_X(\Delta-\tau)}. \end{aligned}$$

The desired discrete-time VAR(1) formulation is obtained by rewriting (I.45):

$$\left[\begin{array}{c} logS_{t+\Delta} - logS_t - r\Delta \\ x_{t+\Delta} \end{array}\right] = \left[\begin{array}{c} (\bar{\mu} - 0.5\sigma^2 - r)\Delta - \frac{\nu}{\kappa_X}\left(1 - e^{-\kappa_X\Delta}\right)\bar{x} \\ \bar{x}\left(1 - e^{-\kappa_X\Delta}\right) \end{array}\right] \tag{I.48}$$
$$+ \left[\begin{array}{cc} 0 & \frac{\nu}{\kappa_X}\left(1 - e^{-\kappa_X\Delta}\right) \\ 0 & e^{-\kappa_X\Delta} \end{array}\right] \cdot \left[\begin{array}{c} logS_t - logS_{t-\Delta} - r\Delta \\ x_t \end{array}\right] + \left[\begin{array}{c} u_{1,t+\Delta} \\ u_{2,t+\Delta} \end{array}\right].$$

Consequently, the equivalence relations between the discrete-time parameters in Equation (35) and the continuous time parameters are:

$$\begin{aligned}
d_1 &= (\bar{\mu} - 0.5\sigma^2 - r)\Delta - \frac{\nu}{\kappa_X}\left(1 - e^{-\kappa_X\Delta}\right)\bar{x}, \\
d_2 &= \bar{x}\left(1 - e^{-\kappa_X\Delta}\right), \\
d_3 &= \frac{\nu}{\kappa_X}\left(1 - e^{-\kappa_X\Delta}\right), \\
d_4 &= e^{-\kappa_X\Delta}, \\
d_5 = \mathrm{Var}[v_{1,t}] &= \left(\sigma^2 + \frac{2\nu\rho\sigma b}{\kappa_X} + \frac{\nu^2 b^2}{\kappa_X^2}\right)\Delta - 2\left(\frac{\nu\rho\sigma b}{\kappa_X^2} + \frac{b^2\nu^2}{\kappa_X^3}\right)\left(1 - e^{-\kappa_X\Delta}\right) + \frac{\nu^2 b^2}{2\kappa^3}\left(1 - e^{-2\kappa_X\Delta}\right), \\
d_6 = \mathrm{Var}[v_{2,t}] &= \frac{b^2}{2\kappa_X}\left(1 - e^{-2\kappa_X\Delta}\right), \\
d_7 = \mathrm{Cov}[v_{1,t}, v_{2,t}] &= \left(\frac{\rho\sigma b}{\kappa_X} + \frac{b^2\nu}{\kappa_X^2}\right)\left(1 - e^{-\kappa_X\Delta}\right) - \frac{\nu b^2}{2\kappa^2}\left(1 - e^{-2\kappa_X\Delta}\right).
\end{aligned}$$

These seven equations yield unique solutions for the continuous time parameters:

$$\begin{aligned}
k_X &= -\log(d_4)\Delta, \\
\bar{x} &= \frac{d_2}{1 - e^{-\kappa_X\Delta}}, \\
\nu &= \frac{d_3\kappa_X}{1 - e^{-\kappa_X\Delta}}, \\
b &= \sqrt{\frac{2d_6\kappa_X}{1 - e^{-2\kappa_X Delta}}}, \\
\rho\sigma b &= \left(d_7 - \left(\frac{b^2\nu}{\kappa_X^2}\left(1 - e^{-\kappa_X\Delta}\right) - \frac{\nu b^2}{2\kappa_X^2}\left(1 - e^{-2\kappa_X\Delta}\right)\right)\right)\left(\frac{\kappa_X}{1 - e^{-\kappa_X\Delta}}\right), \\
\sigma &= \sqrt{\left(d_5 - \left(\frac{2\nu\rho\sigma b}{\kappa_X} + \frac{\nu^2 b^2}{\kappa_X^2}\right)\Delta + 2\left(\frac{\rho\sigma b\nu}{\kappa_X^2} + \frac{b^2\nu^2}{\kappa_X^3}\right)(1 - e^{-\kappa_X\Delta}) - \frac{\nu^2 b^2}{2\kappa_X^3}(1 - e^{-2\kappa_X\Delta})\right)/\Delta}, \\
\alpha &= \frac{d_1 + (0.5\sigma^2 - r)\Delta + (\nu/\kappa_X)(1 - e^{-\kappa_X\Delta})\bar{x}}{\Delta} - \nu\bar{x}.
\end{aligned}$$

Additional Figures and Tables

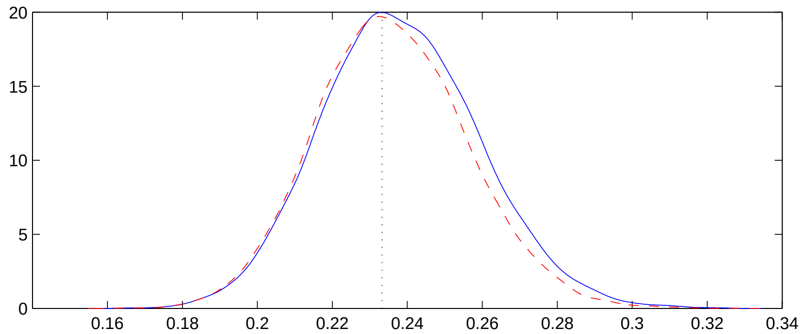

Figure I.1: This figure shows the empirical distributions of point estimates for the predictability coefficient β from Equations (34) and (34). The solid line correspond to classical maximum likelihood estimates, whereas the dashed line depicts the EDF of robust parameter estimates ($c_{Huber} = 5$). The estimates are based on 10,000 simulated sample paths (1,000 observations) which are obtained by simulating returns and log dividend-price ratios according to Equations (34) and (34) with parameters estimated from CRSP data during 1927–1994. In the simulation of observation t, $\omega = -0.40$ is replaced by $\tilde{\omega} = -1.25$ with probability 1/84.

Table I.1: Realized utility in the jump-diffusion case – reversed contamination

	Degree of robustness					
	$\vartheta = 0$	$\vartheta = 0.2$	$\vartheta = 0.6$	$\vartheta = 1.0$	$\vartheta = 2.0$	$\vartheta = 3.0$
Panel (a): In- and out-of-sample JD						
True parameters	1,051	1,051	1,051	1,051	1,051	1,051
(P)MLE	983	1,025	1,035	1,038	1,041	1,041
Robust estimates	985	1,026	1,036	1,038	1,041	1,041
Panel (b): In- and out-of-sample distorted JD						
True parameters	1,169	1,168	1,165	1,164	1,161	1,157
(P)MLE	632	845	959	987	997	1,018
Robust estimates	748	878	980	1,026	1,026	1,030

Notes: The table shows the *ex-post* realized utilities in wealth equivalents of an expected utility maximizer ($\vartheta = 0$) and her robust counterpart. Panel A depicts realized utilities when the data follow a clean jump-diffusion (JD) process in- and out-of-sample for the three cases of knowing the parameters, estimating them with maximum likelihood, and using robust estimates. Panel B shows the same quantities when the data contain small, time-varying distortions to the jump component. Portfolio weights for realized utilities are computed based on five years of daily data generated by the stock process specified in Equation (25) where the time-varying jump size a_t has been reversed relative to the unconditional mean of -1%. The econometric constant used in the robust estimation is chosen to ensure 95% efficiency in case the true data generating process is a clean jump-diffusion ($c_{Huber} = 40$); $\mu = 0.08$, $\sigma = 0.15$, $\lambda_Y = 3$, $\mu_Y = -0.01$, $\sigma_Y = 0.04$.

Table I.2: Realized out-of sample utility

		Classical estimates			Robust estimates		
ϑ		0	1	3	0	1	3
		Dividend yield as predictor					
γ	3	106.76	106.36	105.95	108.06	107.40	106.63
	4	106.37	106.12	105.84	107.45	106.97	106.43
		Random walk					
γ	3	105.83	105.66	105.49	108.02	107.52	106.73
	4	105.66	105.56	105.44	107.52	107.04	106.50
		Mixed strategy (5% conf.)					
γ	3	105.83	105.66	105.49	107.76	107.24	106.53
	4	105.66	105.56	105.44	107.26	106.83	106.34
		Mixed strategy (10% conf.)					
γ	3	105.83	105.66	105.49	108.52	107.91	106.97
	4	105.66	105.56	105.44	107.96	107.37	106.72
		Mixed strategy (20% conf.)					
γ	3	106.29	106.00	105.72	108.66	108.02	107.04
	4	106.01	105.84	105.64	108.07	107.46	106.79
		Mixed strategy (30% conf.)					
γ	3	106.71	106.31	105.92	108.63	108.00	107.03
	4	106.33	106.09	105.82	108.05	107.45	106.77

Notes: The table shows the implied out-of sample utility (wealth equivalent) for an agent using the dividend yield as predictor, the random walk model or a mixed strategy based on evidence for predictability. Utility is reported for different levels of risk aversion (γ) and ambiguity aversion (ϑ) for classical and robust parameter estimates. In each month t, π^* is computed with an investment horizon of $T = 1$ year using parameter estimates based on information up to time t and the current dividend-price ratio. In the mixed strategy cases, the agent tests in each month t for the presence of predictability using the null hypothesis $H_0 : \nu = 0$ against $H_A : \nu > 0$ in Equation (37) at different confidence levels. π^* is then either computed based on a random walk model or based on the optimal policy laid out in Equation (40). The data extends from 1929–2009; the initial sample size for the estimation is 25 years. The econometric constant c_{Huber} used in the estimation equals 6.

Supplemental Appendix B: Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums

Additional Figures and Tables

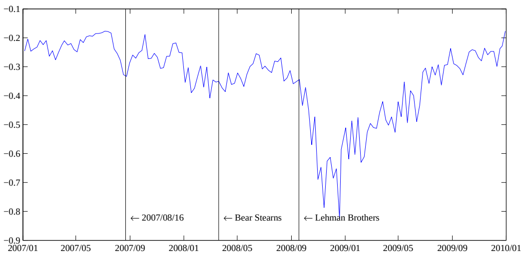

(a) Price impact

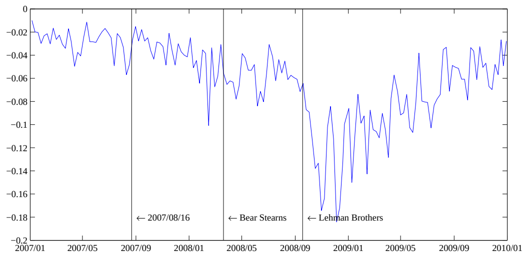

(b) Return reversal

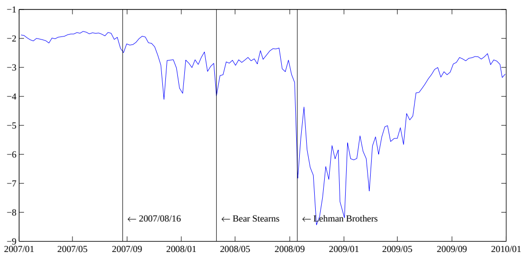

(c) Bid-ask spread

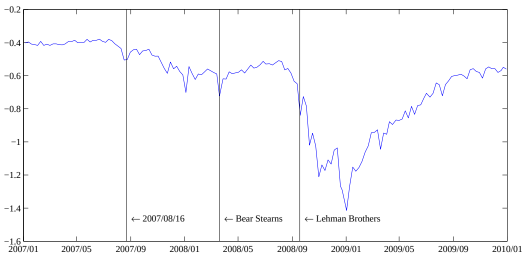

(d) Effective cost

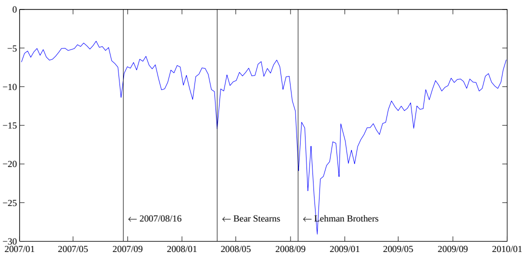

(e) Price dispersion

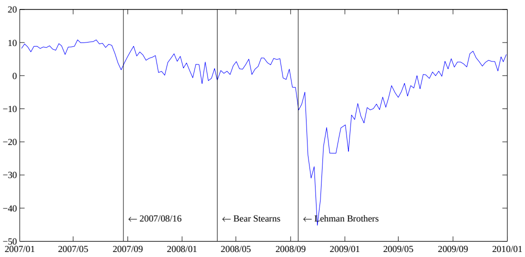

(f) Latent liquidity

Figure I.1: Weekly systematic liquidity. Panels (a)–(e) depict market-wide FX liquidity based on (within measures) averaging of individual exchange rate liquidity (Equation (4)). Latent systematic liquidity obtained from Principle Component Analysis across exchange rates as well as across liquidity measures (Equation (5)) is depicted in Panel (f). The sign of each liquidity measure is adjusted such that the measure represents liquidity rather than illiquidity. The sample is January 2, 2007 – December 30, 2009.

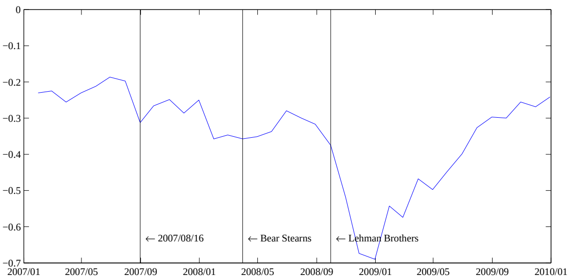

(a) Price impact

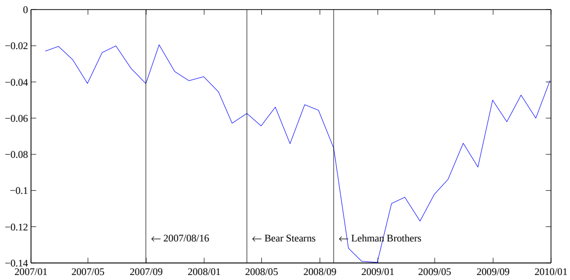

(b) Return reversal

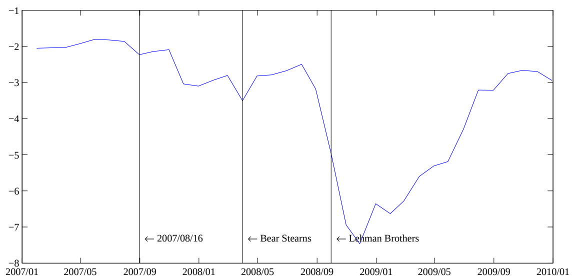

(c) Bid-ask spread

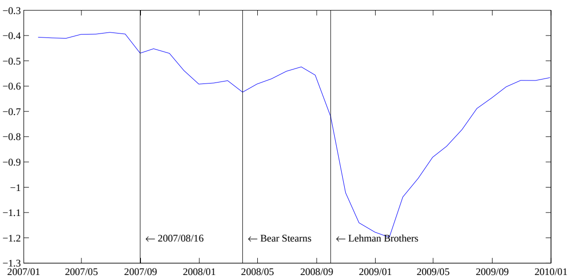

(d) Effective cost

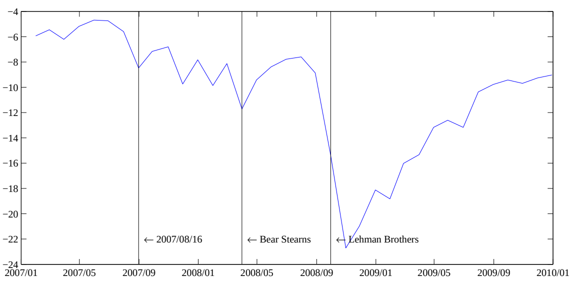

(e) Price dispersion

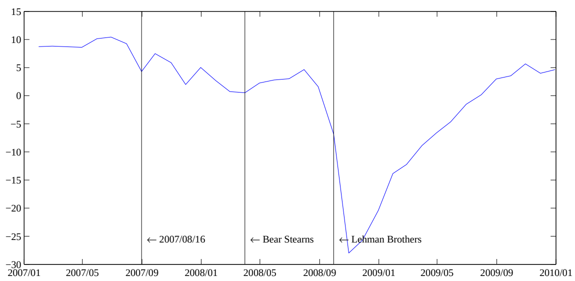

(f) Latent liquidity

Figure I.2: Monthly systematic liquidity. Panels (a)–(e) depict market-wide FX liquidity based on (within measures) averaging of individual exchange rate liquidity (Equation (4)). Latent systematic liquidity obtained from principle component analysis across exchange rates as well as across liquidity measures (Equation (5)) is depicted in Panel (f). The sign of each liquidity measure is adjusted such that the measure represents liquidity rather than illiquidity. The sample is January 2007 – December 2009.

Table I.1: Principle component loadings across exchange rates

	AUD/USD	EUR/CHF	EUR/GBP	EUR/JPY	EUR/USD	GBP/USD	USD/CAD	USD/CHF	USD/JPY
First principle component loadings									
Price impact	0.3026	0.3639	0.2868	0.3756	0.3930	0.3079	0.1718	0.3672	0.3729
Return reversal ($K = 1$)	0.3082	0.3299	0.2893	0.4456	0.3899	0.3672	0.1756	0.1937	0.3984
Return reversal ($K = 3$)	0.3285	0.3487	0.2715	0.3997	0.3816	0.3795	0.1165	0.3011	0.3784
Return reversal ($K = 5$)	0.3129	0.3835	0.2425	0.4049	0.3509	0.3648	0.1333	0.2866	0.4197
Bid-ask spread	0.3109	0.3798	0.2725	0.3801	0.3879	0.2407	0.2200	0.3739	0.3792
Effective cost	0.3270	0.3361	0.3221	0.3426	0.3269	0.3439	0.3247	0.3382	0.3377
Effective cost, volume-weighted	0.3188	0.3378	0.3265	0.3409	0.3335	0.3408	0.3236	0.3411	0.3361
Price dispersion (TSRV, one minute)	0.3310	0.3454	0.3383	0.3550	0.3601	0.2860	0.2932	0.3441	0.3389
Price dispersion (TSRV, five minute)	0.3334	0.3325	0.3398	0.3491	0.3561	0.3185	0.3079	0.3308	0.3294
Average	0.3193	0.3508	0.2988	0.3771	0.3644	0.3277	0.2296	0.3196	0.3656
Second principle component loadings									
Price impact	0.1801	0.1357	0.1733	0.0991	0.0071	−0.0269	−0.9519	−0.0481	−0.0108
Return reversal ($K = 1$)	0.4433	0.2132	0.3078	0.0915	−0.3494	−0.3097	0.2355	−0.6173	−0.0217
Return reversal ($K = 3$)	0.4739	0.2382	0.5237	−0.0072	−0.2554	−0.2057	0.0815	−0.5636	−0.1118
Return reversal ($K = 5$)	0.1526	0.0987	0.5847	−0.1148	−0.3472	0.2776	−0.5598	−0.3174	0.0126
Bid-ask spread	0.1776	0.1114	−0.4483	0.0067	0.0145	−0.6244	0.6005	0.0503	0.0417
Effective cost	−0.5420	−0.2097	0.5770	−0.0866	−0.3448	0.1727	0.3524	0.1946	−0.1048
Effective cost, volume-weighted	0.6606	0.1404	−0.4980	0.1693	0.1181	−0.2247	−0.2972	−0.2651	0.2102
Price dispersion (TSRV, one minute)	−0.0291	−0.0204	−0.1053	0.0646	0.0527	−0.8338	0.4814	0.1146	0.2014
Price dispersion (TSRV, five minute)	−0.4566	−0.0216	0.2978	−0.4013	0.2386	0.2607	0.1529	0.4162	−0.4689
Average	0.1178	0.0762	0.1570	−0.0199	−0.0962	−0.1683	0.0106	−0.1151	−0.0280

Notes: Given a standardized daily measure of liquidity, each row of the table shows principle component loadings for each exchange rate obtained by conducting Principle Component Analysis across the FX rate liquidities. The Principal Component Analysis is repeated for each liquidity measure. The sample is January 2, 2007 – December 30, 2009.

Table I.2: Commonality in liquidity using within measure PCA factors based on FX rates against USD

Measure	*Factor 1*	*Factors 1,2*	*Factors 1,2,3*
	Daily data		
Price impact	0.6280	0.7671	0.8749
Return reversal ($K = 1$)	0.2991	0.4814	0.6335
Return reversal ($K = 3$)	0.3330	0.5049	0.6618
Return reversal ($K = 5$)	0.3205	0.4897	0.6427
Bid-ask spread	0.6682	0.7911	0.9016
Effective cost	0.8831	0.9282	0.9557
Effective cost, volume-weighted	0.8871	0.9321	0.9593
Price dispersion (TSRV, one minute)	0.7729	0.8547	0.9195
Price dispersion (TSRV, five minute)	0.7951	0.8574	0.9152
	Weekly data		
Price impact	0.7231	0.8445	0.9246
Return reversal ($K = 1$)	0.4737	0.6794	0.7908
Return reversal ($K = 3$)	0.5196	0.7165	0.8233
Return reversal ($K = 5$)	0.5302	0.7113	0.8206
Bid-ask spread	0.7751	0.8609	0.9293
Effective cost	0.9085	0.9487	0.9737
Effective cost, volume-weighted	0.9194	0.9577	0.9800
Price dispersion (TSRV, one minute)	0.8644	0.9153	0.9545
Price dispersion (TSRV, five minute)	0.8711	0.9257	0.9602
	Monthly data		
Price impact	0.7951	0.9053	0.9647
Return reversal ($K = 1$)	0.6580	0.8379	0.9123
Return reversal ($K = 3$)	0.6828	0.8357	0.9018
Return reversal ($K = 5$)	0.7184	0.8424	0.9099
Bid-ask spread	0.8604	0.9340	0.9680
Effective cost	0.9227	0.9607	0.9827
Effective cost, volume-weighted	0.9350	0.9676	0.9883
Price dispersion (TSRV, one minute)	0.9118	0.9488	0.9740
Price dispersion (TSRV, five minute)	0.9155	0.9543	0.9806

Notes: For each standardized daily measure of liquidity the first three common factors are extracted using Principle Component Analysis. Then, for each exchange rate and each standardized liquidity measure, liquidity is regressed on its common factors. The table shows the average adjusted-R^2 of these regressions using one, two and three factors. The sample is January 2, 2007 – December 30, 2009. This analysis is conducted using only currency pairs that include the USD.

Table I.3: Principle component loadings across liquidity measures and exchange rates: Average loading for FX rates

	AUD/USD	EUR/CHF	EUR/GBP	EUR/JPY	EUR/USD	GBP/USD	USD/CAD	USD/CHF	USD/JPY
					Daily data				
PC1	0.3085	0.0444	0.1021	0.0790	0.0297	0.1707	0.1699	0.0453	0.0374
PC2	−0.1442	0.0087	0.0619	0.0192	0.0116	0.1679	0.1162	0.0219	0.0087
PC3	0.1335	0.0070	−0.0141	0.0112	0.0008	−0.0535	−0.1854	−0.0018	0.0044
					Weekly data				
PC1	0.3123	0.0477	0.1076	0.0865	0.0320	0.1735	0.1574	0.0481	0.0395
PC2	−0.1791	0.0147	0.1189	0.0304	0.0208	0.2414	0.0884	0.0371	0.0096
PC3	0.1131	0.0050	−0.0282	0.0150	0.0005	0.0009	−0.1816	−0.0062	0.0034
					Monthly data				
PC1	0.3079	0.0505	0.1139	0.0914	0.0342	0.1859	0.1543	0.0510	0.0411
PC2	−0.2223	0.0024	0.1298	0.0103	0.0164	0.2270	0.1160	0.0327	0.0026
PC3	0.0295	0.0013	0.0244	−0.0009	0.0019	0.1033	−0.2001	−0.0090	−0.0100

Notes: Principle component loadings across FX liquidity measures and exchange rates are extracted by Principle Component Analysis. The table reports the average loading for each exchange rate at different time frequencies. The sample is January 2, 2007 – December 30, 2009.

Table I.4: Principle component loadings across liquidity measures and exchange rates: Average loading for liquidity measures

	Return reversal	Price impact	Bid-ask spread	Effective cost	Price dispersion
			Daily data		
PC1	0.1182	0.1161	0.1170	0.1025	0.0944
PC2	−0.0666	0.0430	0.1027	0.0439	0.0281
PC3	−0.0988	0.0418	−0.0093	0.0003	0.0115
			Weekly data		
PC1	0.1219	0.1217	0.1087	0.1063	0.0995
PC2	−0.0236	0.0302	0.1034	0.0705	0.0318
PC3	−0.1045	0.0558	−0.0097	−0.0014	0.0166
			Monthly data		
PC1	0.1185	0.1260	0.1100	0.1129	0.1049
PC2	0.0275	−0.0037	0.0869	0.0540	0.0102
PC3	−0.1037	0.0256	0.0474	−0.0049	0.0024

Notes: Principle component loadings across FX liquidity measures and exchange rates are extracted by Principle Component Analysis. The table reports the average loading for each measure of liquidity at different time frequencies. The sample is January 2, 2007 – December 30, 2009.

Table I.5: Further evidence for commonality

Liquidity measure	Mean β	Std. β	% pos.	% pos. & signif.	Adj.-R^2
			Daily data		
Price impact	0.3405	0.1054	100.00%	77.78%	0.0143
Return reversal ($K = 1$)	0.0169	0.1462	55.56%	22.22%	0.0048
Return reversal ($K = 3$)	0.1401	0.1409	77.78%	55.56%	0.0059
Return reversal ($K = 5$)	0.1613	0.1417	77.78%	44.44%	0.0073
Bid-ask spread	0.4855	0.1899	88.89%	88.89%	0.1062
Effective cost	0.9460	0.0614	100.00%	100.00%	0.3461
Effective cost, volume-weighted	0.8842	0.0660	100.00%	100.00%	0.2955
Price dispersion (TSRV, one minute)	0.8780	0.0665	100.00%	100.00%	0.3739
Price dispersion (TSRV, five minute)	0.8873	0.0534	100.00%	100.00%	0.3793
			Weekly data		
Price impact	0.6932	0.2227	100.00%	77.78%	0.0723
Return reversal ($K = 1$)	0.6971	0.2964	88.89%	44.44%	0.0437
Return reversal ($K = 3$)	0.4502	0.2793	100.00%	44.44%	0.0588
Return reversal ($K = 5$)	0.1971	0.2705	88.89%	55.56%	0.0572
Bid-ask spread	0.5615	0.2689	88.89%	77.78%	0.3465
Effective cost	1.0746	0.1027	100.00%	100.00%	0.4871
Effective cost, volume-weighted	1.0635	0.1048	100.00%	100.00%	0.4622
Price dispersion (TSRV, one minute)	1.0833	0.0929	100.00%	100.00%	0.5648
Price dispersion (TSRV, five minute)	1.0268	0.0810	100.00%	100.00%	0.5639
			Monthly data		
Price impact	1.2571	0.3650	100.00%	88.89%	0.3887
Return reversal ($K = 1$)	0.8301	0.5774	88.89%	44.44%	0.1623
Return reversal ($K = 3$)	0.8909	0.5750	88.89%	44.44%	0.1209
Return reversal ($K = 5$)	0.8105	0.5487	88.89%	55.56%	0.1574
Bid-ask spread	1.1452	0.2941	100.00%	88.89%	0.5480
Effective cost	1.3638	0.1759	100.00%	100.00%	0.7129
Effective cost, volume-weighted	1.3332	0.1798	100.00%	100.00%	0.7129
Price dispersion (TSRV, one minute)	1.1339	0.1462	100.00%	100.00%	0.7007
Price dispersion (TSRV, five minute)	1.1136	0.1367	100.00%	100.00%	0.7047

Notes: This table shows time series regression results when daily relative changes in individual exchange rate j liquidity are regressed on relative changes in systematic FX liquidity. The latter is given by the average liquidity across exchange rates, without exchange rate j, similarly to Chordia, Roll, and Subrahmanyam (2000). Mean β and Std. β denote cross-sectional average and standard deviation of slope coefficients. % pos. and % pos. & signif. denote the percentages of estimates which are positive as well as positive and significantly different from zero. The last column shows the adjusted-R^2. The sample is January 2, 2007 – December 30, 2009.

Table I.6: Further evidence for liquidity spirals in the FX market

	$const$	$L^{pca}_{M,t-1}$	VIX_{t-1}	$LIBOIS_{t-1}$	Adj. R^2
Coefficient	16.7529		−0.5270	−4.9014	0.7877
Std. error	(0.8049)		(0.0416)	(0.9838)	
Coefficient	8.4177			−13.3316	0.6353
Std. error	(0.7039)			(1.4024)	
Coefficient	10.6614	0.3553	−0.3294	−3.3678	0.8135
Std. error	(1.2812)	(0.0762)	(0.0491)	(0.6365)	
Coefficient	3.2151	0.6234		−5.0928	0.7503
Std. error	(0.4617)	(0.0465)		(0.8589)	

Notes: Regression of daily latent systematic FX liquidity ($L^{(pca)}_{M,t}$) on lagged VIX and LIBOR-OIS spread. Four different specifications of the regression model are estimated. Heteroscedasticity and autocorrelation (HAC) robust standard errors are shown in parenthesis. The sample is January 2, 2007 – December 30, 2009.

Supplemental Appendix C: The Joint Dynamics of Hedge Fund Returns, Liquidity, and Volatility

Description of Hedge Fund Styles

Convertible Arbitrage

Convertible arbitrage hedge funds try to exploit pricing anomalies between convertible bonds and the corresponding equity. Frequently, the option embedded in convertible bonds is undervalued with respect to the theoretical value given the price of the underlying stock. Thus, a typical investment strategy is to buy the convertible bond and hedge the equity risk by selling short the stock.

Dedicated Short Bias

In order to be included in this category, hedge fund managers must constantly have a short bias in their position. As Dedicated short hedge funds use predominantly short positions, they can be viewed as mirrors of traditional long only managers.

Emerging Markets

Emerging market hedge funds invest in various types of securities in emerging markets. As these markets are usually not as liquid and less frequently covered by analysts than developed countries, hedge funds which are specialized in these emerging economies try to profit by investing in undervalued securities.

Equity Market Neutral

Equity market neutral funds rely on quantitative portfolio construction techniques to exploit equity market inefficiencies between related equity securities. The strategy seeks to construct portfolios in such a way that the fund is not exposed to market risk.

Event Driven

As the name suggests, this strategy focuses on investing in securities issued by companies for which a corporate event such as a spin-off, merger, acquisition, bankruptcy reorganization or re-capitalization is anticipated. Different subcategories are summarized by this style, namely, distressed securities, risk or merger arbitrage and multi strategy.

Distressed

Distressed securities funds invest in equity, debt or trade claims of companies that are in financial distress or bankruptcy. These securities typically trade at substantial discounts that hedge funds try to exploit when they perceive a turnaround.

Risk Arbitrage

Hedge funds following this strategy simultaneously investment in both companies involved in a merger or acquisition. In a typical trade the fund is long the stock of the company being acquired and short the stock of the acquirer.

Multi-Strategy

This strategy involves investing in various event driven strategies in response to market opportunities. In addition to the above subcategories, a popular strategy is to invest in Regulation D companies, i.e. micro and small capitalization companies that are raising money in private capital markets.

Fixed Income Arbitrage

This strategy seeks to exploit from price anomalies between interest rate securities in global fixed income markets.

Global Macro

Global macro hedge funds have the broadest investment universe and invest in any of the world's capital or derivative markets on a discretionary basis.

Long/Short Equity

This strategy involves equity investing on both the long and short sides without necessarily being market neutral. Most long/short equity hedge funds maintain a net long position, thus, their performance usually correlates with traditional mutual funds.

Managed Futures

Managed futures managers, which are usually referred to as Commodity Trading Advisors (CTAs), trade listed financial and commodity futures contracts either on a systematic or an discretionary basis.

Additional Figures and Tables

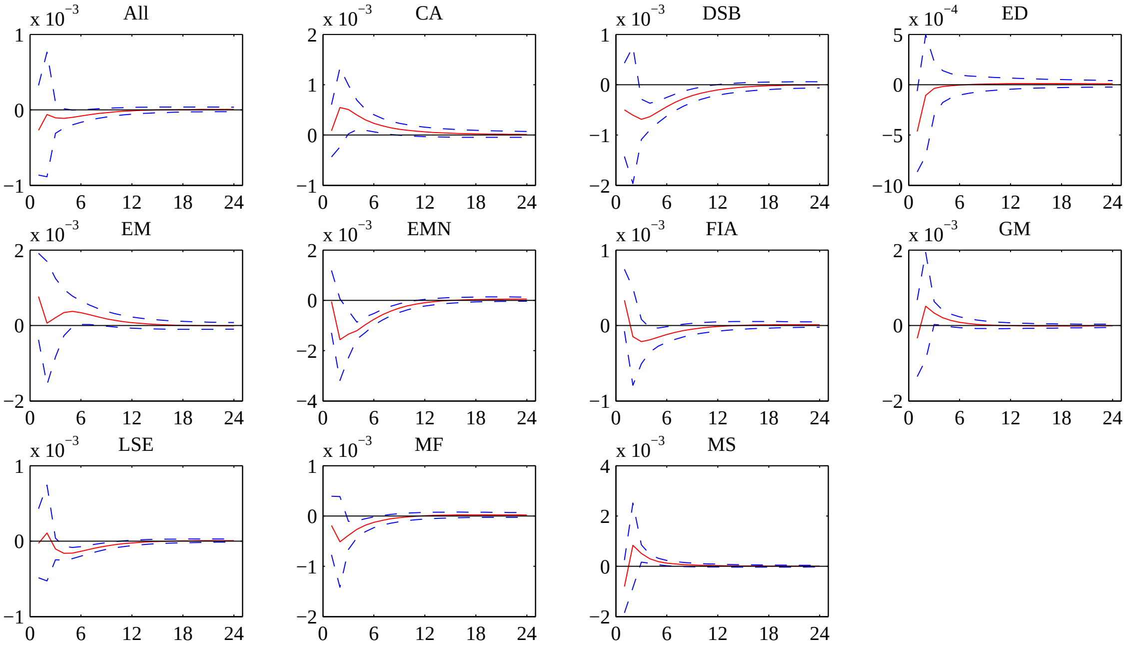

(a) Response to shock in VOL_{EQ}

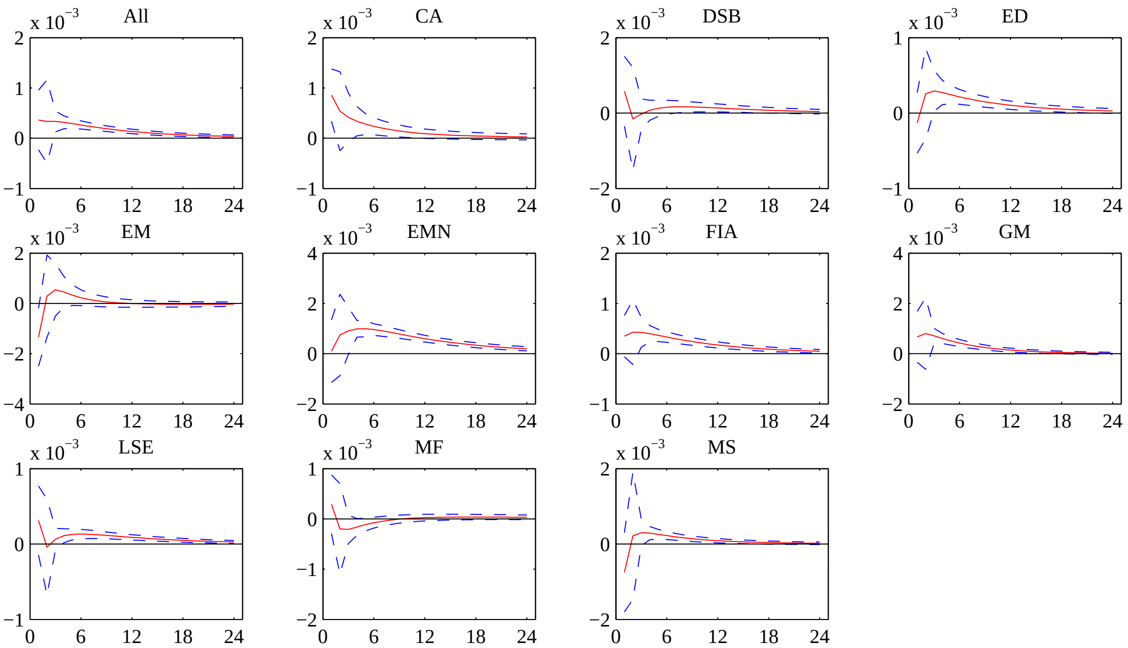

(b) Response to shock in ILL_{EQ}

Figure I.1: Impulse response functions of various hedge fund returns to shocks in equity and FX volatility and illiquidity for different ordering of endogenous variables.

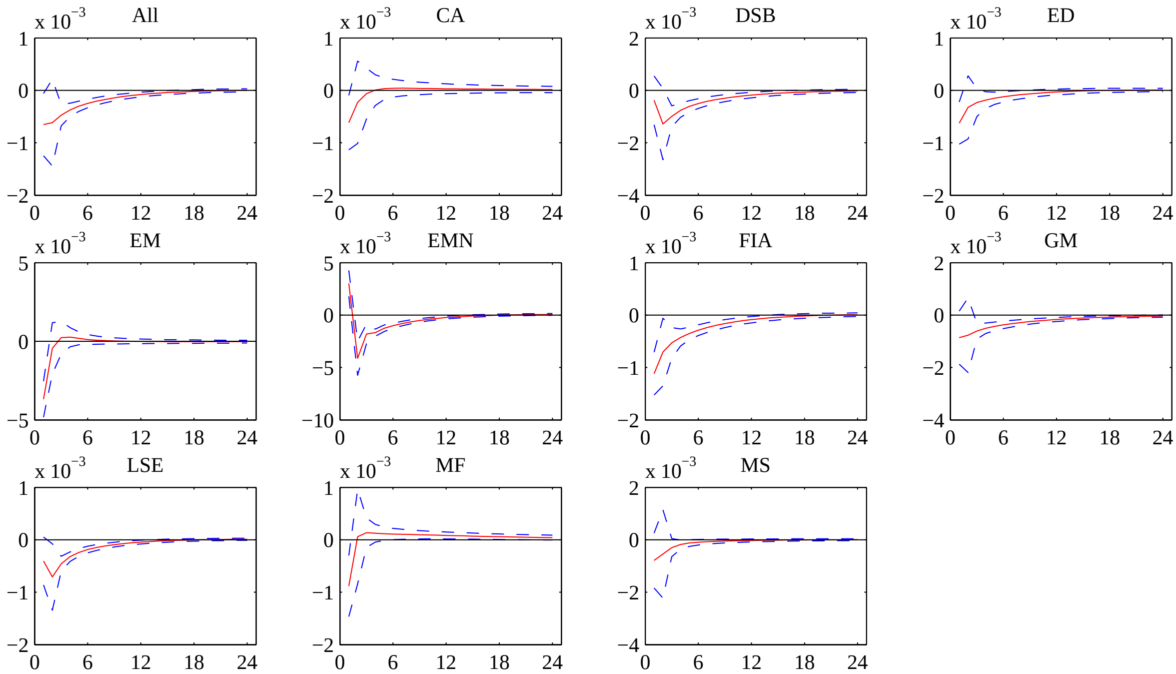

(c) Response to shock in VOL_{FX}

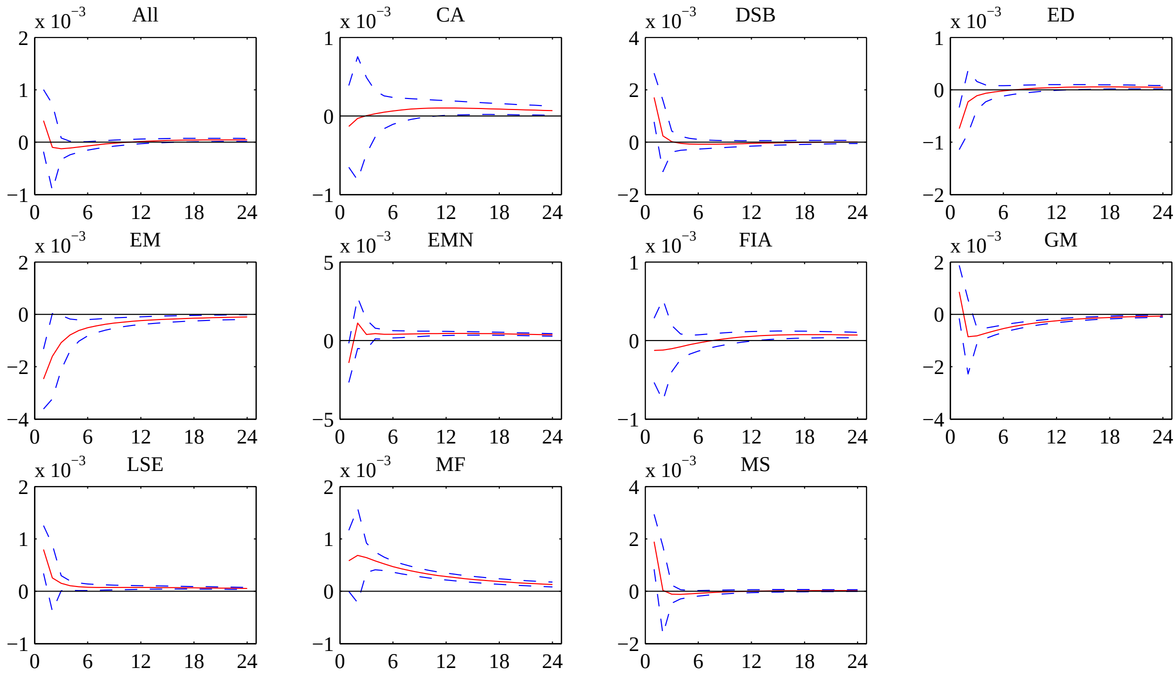

(d) Response to shock in ILL_{FX}

Figure I.1 (continued): Impulse response functions of various hedge fund returns to shocks in equity and FX volatility and illiquidity for different ordering of endogenous variables.

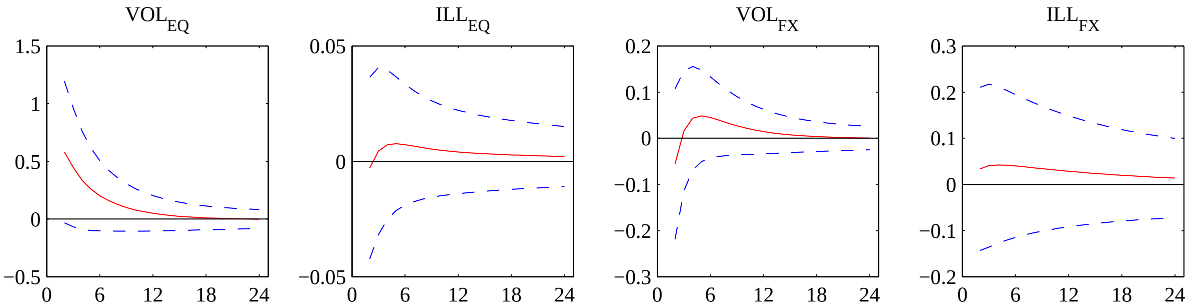

(a) Response to shock in $r_{HF,all}$

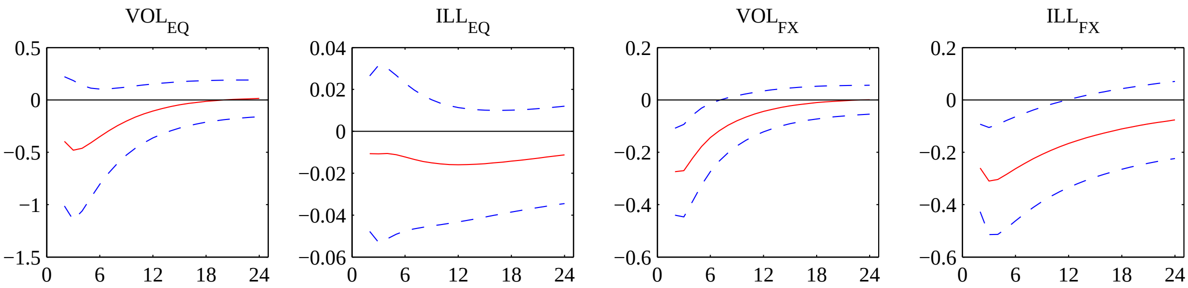

(b) Response to shock in $r_{HF,ca}$

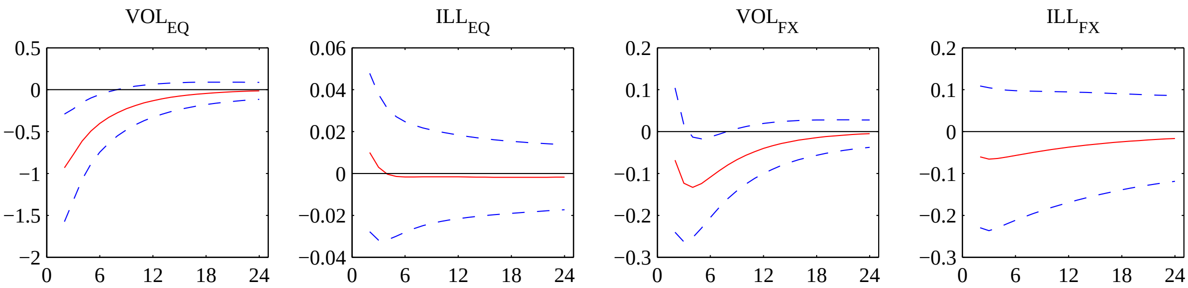

(c) Response to shock in $r_{HF,dsb}$

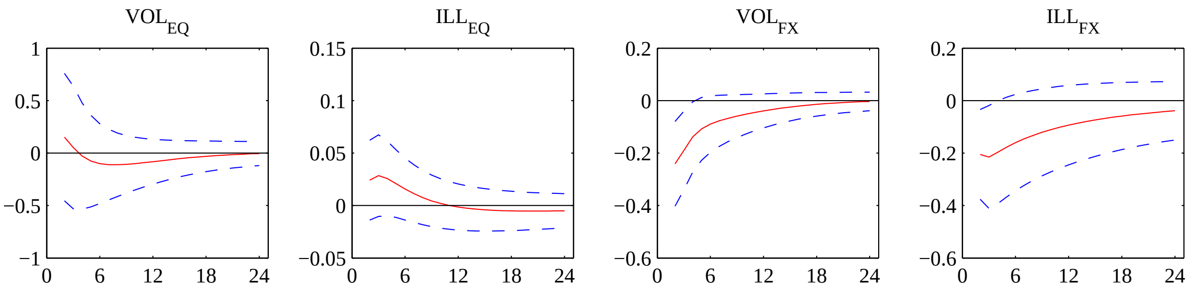

(d) Response to shock in $r_{HF,ed}$

Figure I.2: Impulse response functions of equity and foreign exchange illiquidity and volatility for different ordering of endogenous variables.

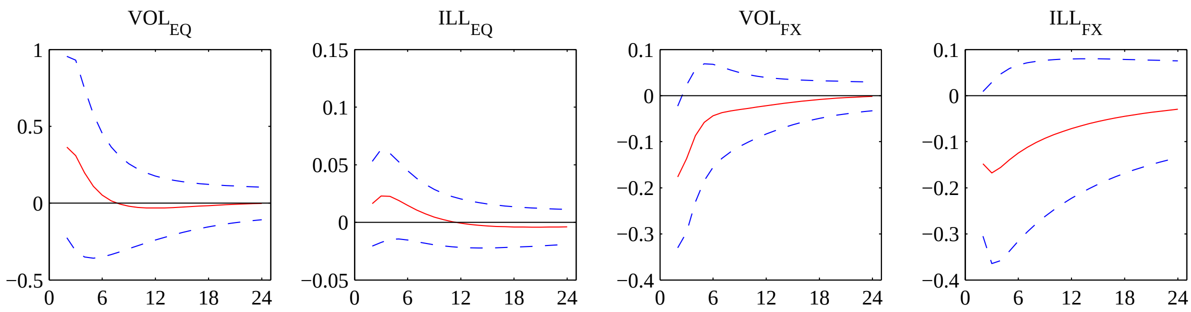

(e) Response to shock in $r_{HF,em}$

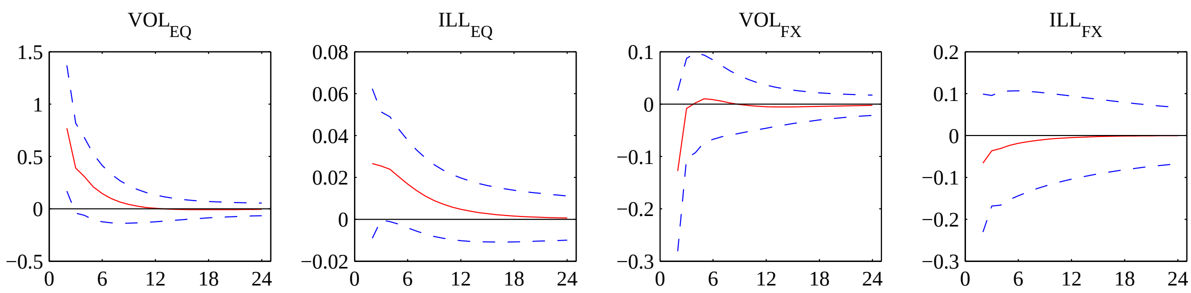

(f) Response to shock in $r_{HF,emn}$

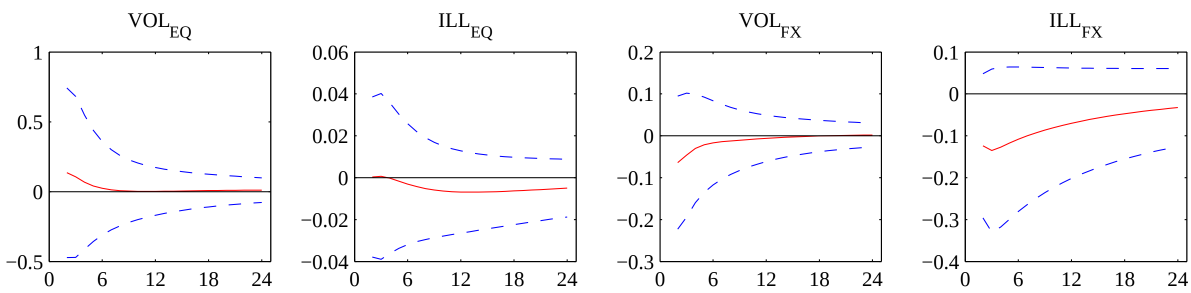

(g) Response to shock in $r_{HF,fia}$

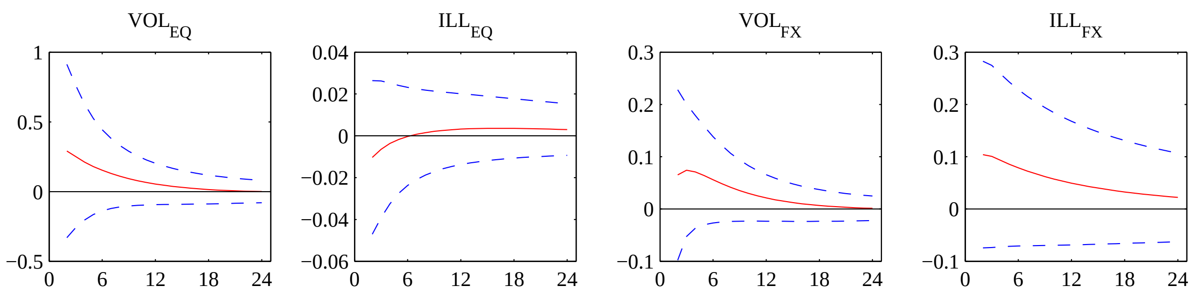

(h) Response to shock in $r_{HF,gm}$

Figure I.2 (continued): Impulse response functions of equity and foreign exchange illiquidity and volatility for different ordering of endogenous variables.

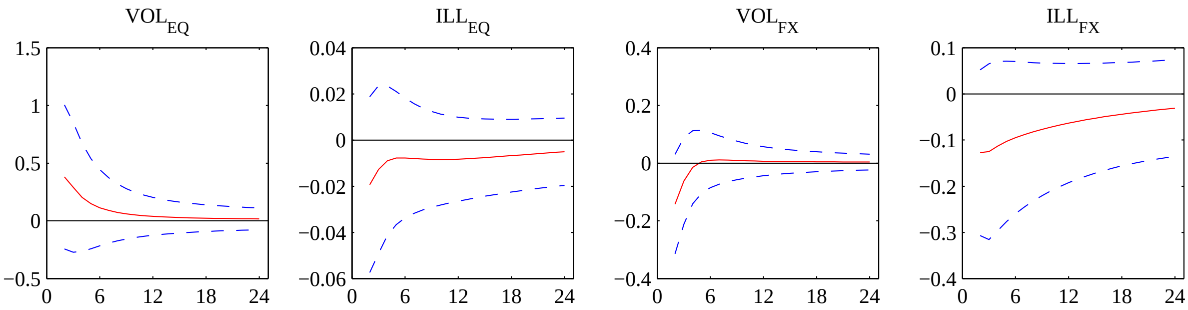

(i) Response to shock in $r_{HF,lse}$

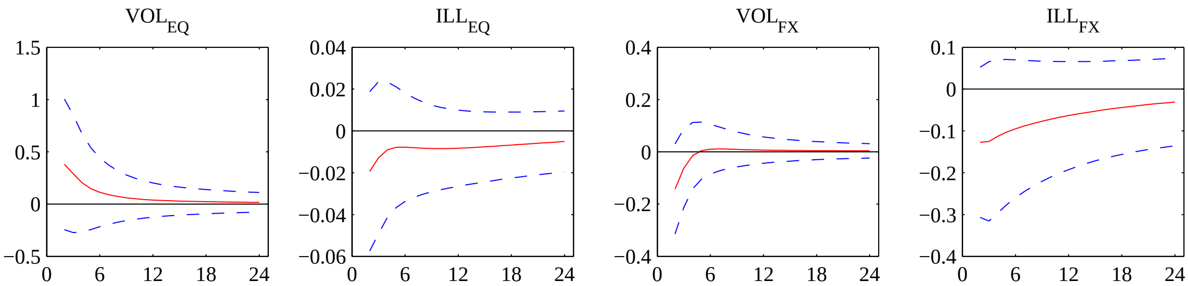

(j) Response to shock in $r_{HF,mf}$

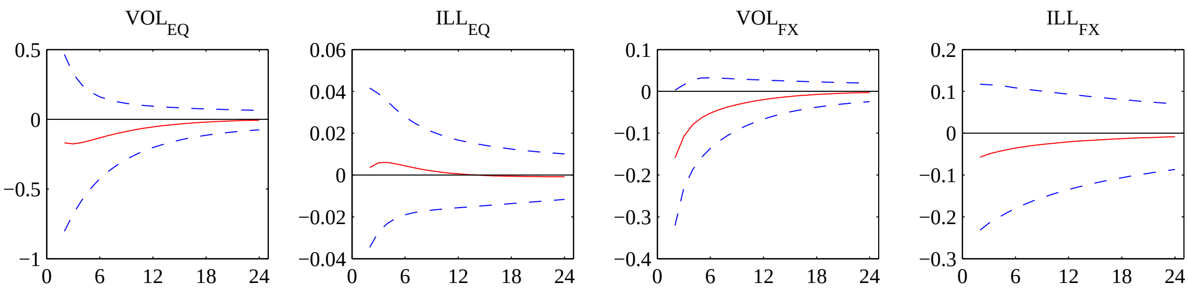

(k) Response to shock in $r_{HF,ms}$

Figure I.2 (continued): Impulse response functions of equity and foreign exchange illiquidity and volatility for different ordering of endogenous variables.

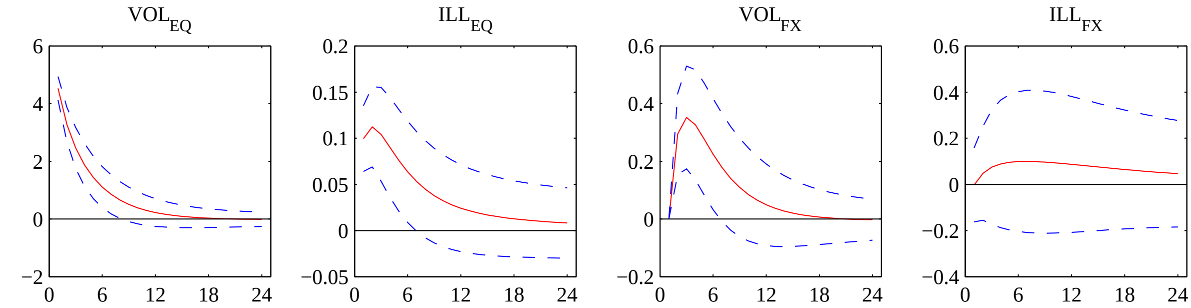

(a) Response to shock in VOL_{EQ}

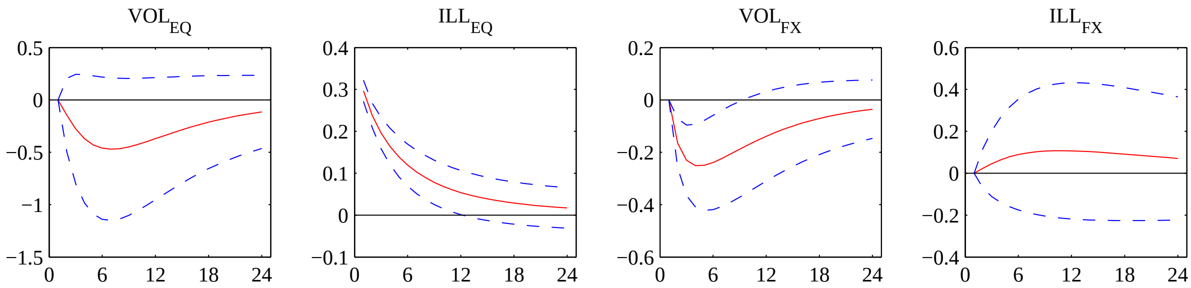

(b) Response to shock in ILL_{EQ}

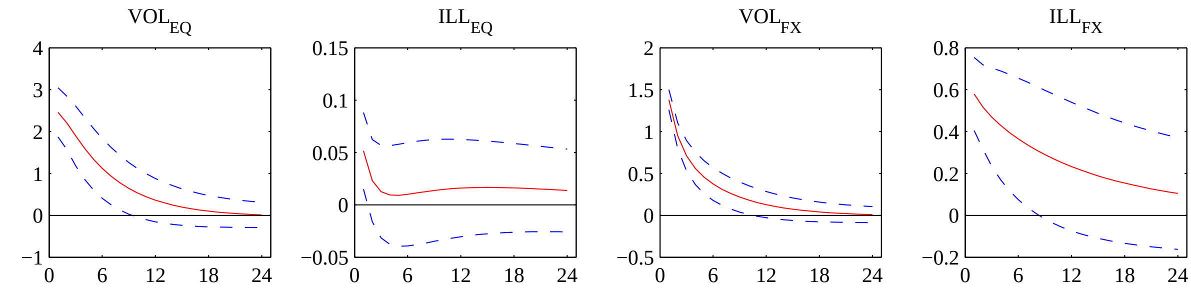

(c) Response to shock in VOL_{FX}

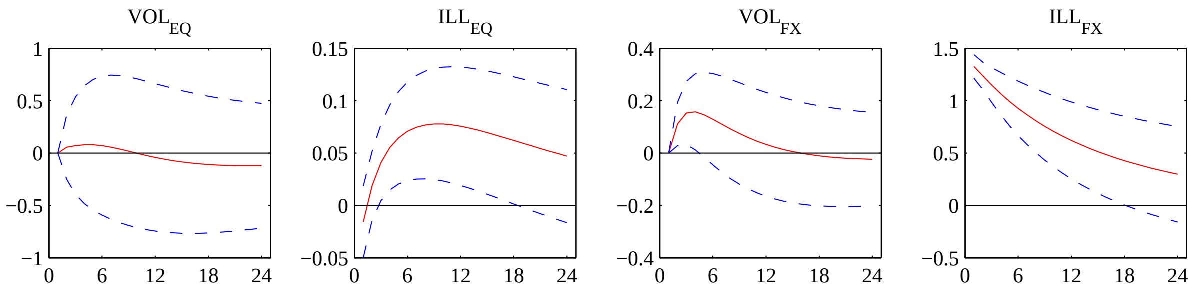

(d) Response to shock in ILL_{FX}

Figure I.3: Impulse response functions of equity and foreign exchange illiquidity and volatility for different ordering of endogenous variables.

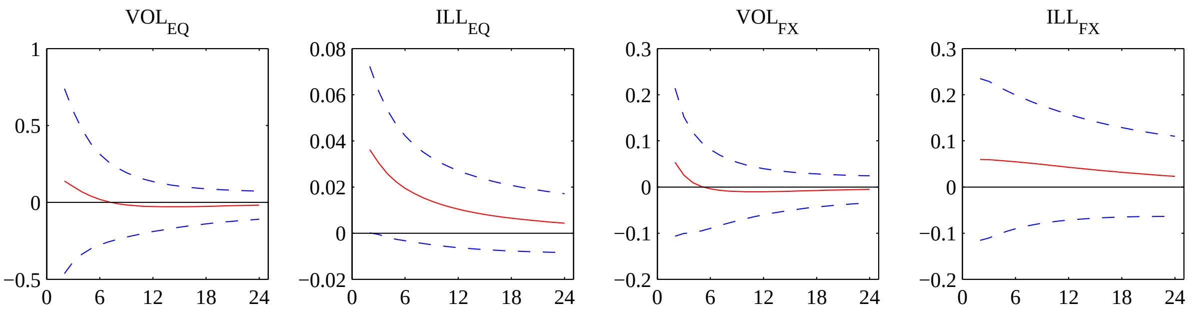

(a) Response to shock in $PTFSBD$

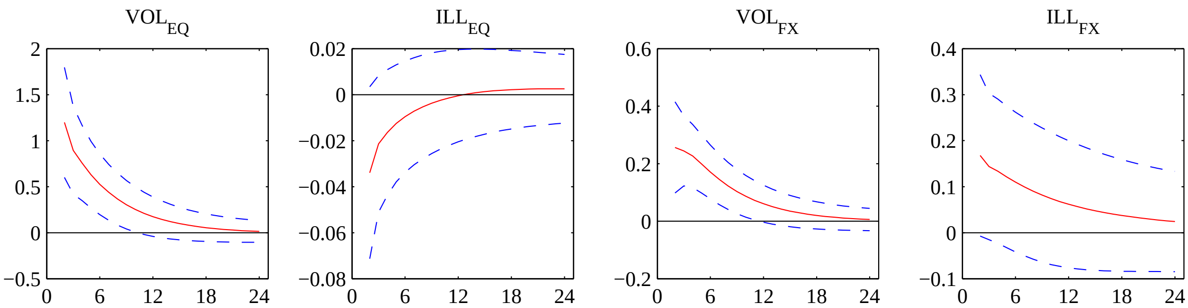

(b) Response to shock in $PTFSCOM$

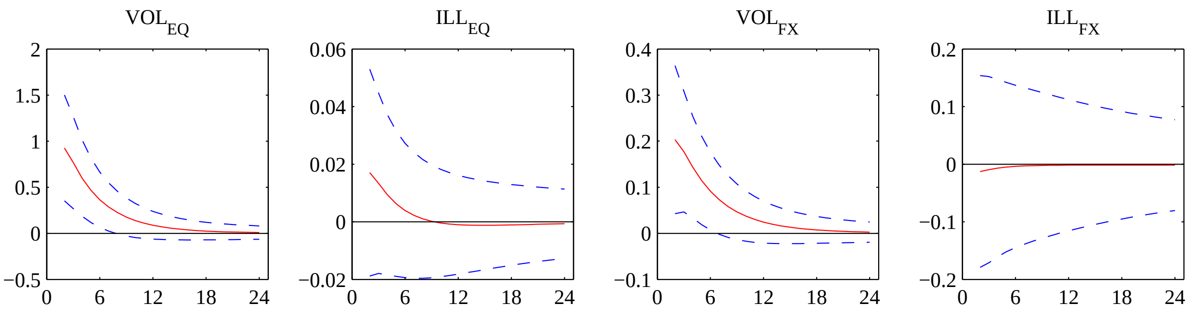

(c) Response to shock in $PTFSFX$

Figure I.4: Impulse response functions of equity and foreign exchange illiquidity and volatility to trend following risk factors for different ordering of endogenous variables.

Table I.1: This table shows Fung-Hsieh factor loadings (β's) of various hedge fund styles (standard errors in parenthesis).

	$r_{HF,all}$	$r_{HF,ca}$	$r_{HF,dsb}$	$r_{HF,ed}$	$r_{HF,em}$	$r_{HF,emn}$	$r_{HF,fia}$	$r_{HF,gm}$	$r_{HF,lse}$	$r_{HF,mf}$	$r_{HF,ms}$
α	0.0025***	0.0025***	0.0010*	0.0029***	0.0013*	0.0014*	0.0016***	0.0038***	0.0025***	0.0024***	0.0026***
	(0.0005)	(0.0005)	(0.0008)	(0.0004)	(0.0011)	(0.0012)	(0.0004)	(0.0009)	(0.0006)	(0.0010)	(0.0004)
$MKT - RF$	0.1065***	0.0363***	−0.3522***	0.0869***	0.2068***	0.0701***	0.0278***	0.0590***	0.1742***	−0.0148	0.0321***
	(0.0120)	(0.0118)	(0.0195)	(0.0087)	(0.0258)	(0.0274)	(0.0093)	(0.0207)	(0.0128)	(0.0229)	(0.0097)
SMB	0.0372***	−0.0034	−0.1478***	0.0240***	0.0817***	−0.0074	−0.0167*	−0.0027	0.1007***	−0.0100	0.0004
	(0.0142)	(0.0140)	(0.0232)	(0.0104)	(0.0307)	(0.0326)	(0.0110)	(0.0246)	(0.0152)	(0.0272)	(0.0116)
$TERM$	−0.0048**	−0.0101***	−0.0006	−0.0011	−0.0027	−0.0040	−0.0080***	−0.0079**	−0.0043*	−0.0062*	−0.0049***
	(0.0025)	(0.0025)	(0.0041)	(0.0018)	(0.0054)	(0.0057)	(0.0019)	(0.0043)	(0.0027)	(0.0048)	(0.0021)
$CREDIT$	−0.0136***	−0.0283***	−0.0126***	−0.0107***	−0.0145***	−0.0187***	−0.0273***	−0.0147***	−0.0051*	−0.0038	−0.0184***
	(0.0031)	(0.0031)	(0.0051)	(0.0023)	(0.0067)	(0.0071)	(0.0024)	(0.0054)	(0.0033)	(0.0060)	(0.0025)
$PTFSBD$	−0.0114***	−0.0056*	0.0023	−0.0120***	−0.0200***	−0.0189***	−0.0049**	−0.0101*	−0.0081***	0.0163***	−0.0030
	(0.0036)	(0.0036)	(0.0059)	(0.0027)	(0.0079)	(0.0083)	(0.0028)	(0.0063)	(0.0039)	(0.0070)	(0.0031)
$PTFSFX$	0.0052**	−0.0026	−0.0012	0.0028*	−0.0028	0.0112**	−0.0042**	0.0062*	0.0033	0.0182***	0.0032*
	(0.0029)	(0.0029)	(0.0047)	(0.0021)	(0.0063)	(0.0067)	(0.0023)	(0.0050)	(0.0031)	(0.0055)	(0.0024)
$PTFSCOM$	0.0079**	−0.0035	−0.0086*	0.0024	0.0076	0.0114*	0.0024	0.0090*	0.0050*	0.0180***	0.0002
	(0.0040)	(0.0040)	(0.0066)	(0.0029)	(0.0087)	(0.0092)	(0.0031)	(0.0070)	(0.0043)	(0.0077)	(0.0033)

***, **, and * indicate significance at the 5%, 10%, and 25% levels, respectively.

Jan Wrampelmeyer - Curriculum Vitae

Educational Background

09/2007 – 04/2011	**University of Zurich, Switzerland** Swiss Finance Insitute PhD Program in Finance
09/2006 – 08/2007	**Maastricht University, The Netherlands** Master of Science in Econometrics and Operations Research
09/2003 – 08/2006	**Maastricht University, The Netherlands** Bachelor of Science in Econometrics and Operations Research

Academic Exchanges and Visiting Positions

01/2010 – 06/2010	**Columbia Business School, Columbia University, New York, USA** Chazen Visiting Scholar
08/2005 – 12/2005	**W.P. Carey School of Business, Arizona State University, Tempe, USA** Exchange Student

Work Experience

09/2008 – 04/2011	**University of Zurich, Switzerland** Research Assistant at the Department of Banking and Finance
09/2009 – 12/2009	**Swiss Federal Institute of Technology Lausanne, Switzerland** Teaching Assistant
06/2006 – 08/2006	**HSBC Trinkaus & Burkhardt AG, Düsseldorf, Germany** Intern at the Equity Derivatives Group

Professional Qualifications

CAIA	Passed Level II of the Chartered Alternative Investment Analyst examination in March 2010

Awards and Scholarships

04/2010	Eastern Finance Association 2010 Outstanding International Paper Award
01/2010 – 06/2010	Swiss National Science Foundation fellowship for prospective researchers
09/2007 – 09/2008	Grant provided by Swiss Finance Institute during the first year of doctoral studies
08/2006	Beta Gamma Sigma Honor Society lifelong membership award
12/2005	Dean's Honors List, W.P. Carey School of Business, Arizona State University