Essays on Asset Pricing

Tuomas Tomunen

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2020

Abstract

Essays on Asset Pricing

Tuomas Tomunen

How are the prices of financial assets determined? In this dissertation, I test various theories empirically, focusing on several classes of bonds. In the first chapter, I test whether asset prices reflect the risk-exposures of financial intermediaries in a setting that is well suited to tackling concerns about omitted risk factors. I analyze catastrophe bonds whose cash flows are linked to the occurrence of natural disasters and find that 71% of the variation in their expected returns can be explained by a theoretically-motivated measure of financial intermediaries' marginal rate of substitution. Assuming that natural disasters are independent of aggregate wealth, this pricing result is inconsistent with any explanation based on macroeconomic risk factors. However, the result is consistent with intermediary asset pricing models that suggest that financial intermediaries are marginal investors in capital markets. I also show that the premium on natural disaster risk has decreased significantly in recent years and has become less responsive to the occurrence of disasters, suggesting that intermediaries' access to outside capital has improved over time. In the second chapter, which is coauthored with Robert J. Hodrick, we examine the statistical term structure model of Cochrane and Piazzesi (2005) and its affine counterpart, developed in Cochrane and Piazzesi (2008), in several out-of-sample analyzes. The model's one-factor forecasting structure across bonds with two, three, four, and five years to maturity characterizes the term structures of additional major currencies in samples ending in 2003. In post-2003 data such one-factor structures again characterize each currency's term structure, but we reject equality of the coefficients across the two samples. We derive currency return forecasting implications from the Cochrane and Piazzesi (2008) affine model showing that the term structure forecasting variables in each currency should predict cross-currency investments, but we find no support for these predictions in either pre-2004 or post-2003 data, whereas the interest differentials do predict currency returns. Here too, though, we find strong evidence of parameter instability as the parameter estimates on the interest differentials change sign. In recursive out-of-sample forecasts of excess rates of return on bonds in each currency, the Cochrane and Piazzesi (2008) term structure forecasting models fail to beat forecasts from the historical average excess rates of return. Graphical analysis indicates that the instability in the forecasting models' parameters begins in the global financial crisis.

Table of Contents

List of Tables

List of Figures

Acknowledgements

I wish to express my deepest gratitude to my advisors Paul Tetlock, Robert Hodrick, Kent Daniel, and Suresh Sundaresan for guidance and feedback that has made my studies a thoughtful and rewarding journey. I am grateful to my coauthors and collaborators Petri Jylhä, Matti Suominen, Lars Lochstoer, and Renxuan Wang for all the discussions and meetings that have taught me a lot. I am also grateful for Matti Keloharju and Wei Jiang for invaluable early mentoring that helped to shape my thinking.

I am thankful to Adrien, Aref, Yifeng, Lira, Tomas, and Daheng for continuous support during our studies and for friendship outside the classroom.

Finally, I am indebted to my family. To my father, Mauri, for sparking my interest towards stock markets at an early age. To my mother, Marjo-Tea, for all the support, especially during the last year of my studies. To my brother and sister, Lauri and Hanna, for all the encouragement. To my wife, Mari, for embarking on this exciting journey with me, and to my son Hugo, for joining it a year ago.

Dedication

To Mari and Hugo.

Chapter 1: Failure to Share Natural Disaster Risk

Capital markets improve welfare by enabling risk sharing across members of the society. However, recent research recognizes that frictions in financial intermediation, such as limited investor knowledge and intermediary capital, could disrupt risk sharing across investors (He and Krishnamurthy, 2013; Brunnermeier and Sannikov, 2014). A key implication of these frictions is that asset prices will reflect the preferences of intermediaries, such as market specialists, rather than those of investors. Several recent studies including Adrian, Etula, and Muir (2014) empirically test this intermediary approach to asset pricing. Most tests specify an empirical proxy for intermediaries' marginal rate of substitution ("intermediary SDF") and show that it can price a broad cross-section of assets.

While such tests provide strong *prima facie* evidence that financial intermediaries are marginal investors, they are subject to the omitted variable criticism: the chosen proxy for the intermediary SDF could be correlated with unobserved risk factors that can arise even if intermediaries are not marginal (e.g., Santos and Veronesi, 2018). To overcome this problem, one would need an instrument that is highly correlated with the marginal utility of intermediaries while being uncorrelated with the marginal utility of other agents in the economy. However, because marginal utility is not directly observable and because macroeconomic measures are highly interconnected, finding a variable that satisfies this exclusion restriction is not straightforward.

In this paper, I test these theories of risk sharing in a market that is uniquely well suited to address this identification challenge. In particular, I analyze catastrophe bonds ("cat bonds"), whose cash flows are linked to the occurrence of natural disasters. These securities have little exposure to traditional macroeconomic risks[1] and are mainly held by specialized asset managers

[1] Cat bonds are intentionally structured in a way that minimizes their exposure to traditional fixed income risks, such as interest rate risk (cat bonds generally pay floating rate coupons) or credit risk (cat bonds are fully collateralized), to provide investors with "pure" exposure to natural disaster risk.

who are potentially marginal investors in this market.

My main finding is that 71% of the variation in the expected returns of the test assets is explained by a theoretically-motivated measure of these intermediaries' marginal rate of substitution. Figure 1.1 illustrates this result by plotting the expected excess returns on individual cat bonds against the predicted values from the one-factor intermediary model. The factor consists of asset class specific risk linked to the occurrence of natural disasters. Under the identifying assumption that natural disasters don't have a first-order causal effect on the marginal value of aggregate wealth, natural disaster risk is diversifiable and won't command a risk premium if prices reflect the preferences of well-diversified outside investors—regardless of their exact preferences.[2] A financial intermediary, on the other hand, who specializes in natural disaster risk is exposed to the asset class specific risk and might demand a premium for holding it if he cannot pass the risk efficiently to the outside investors.

Interpreting the pattern in Figure 1.1 as evidence for financial intermediaries being marginal investors relies on three assumptions. First, catastrophe risk must be diversifiable in equilibrium. While the sheer difference in magnitude between the worst-case catastrophe losses at the macroeconomic-level (estimated at hundreds of billions of dollars) and the value of global financial assets (hundreds of trillions of dollars) suggests that this is likely to be the case, I provide empirical evidence to support the plausibility of this assumption. Using 66 years of macroeconomic and natural disaster data on 13 developed countries, I cannot reject the null hypothesis that natural disasters do no systematically affect consumption or GDP growth, house prices, or stock market returns. Furthermore, in a shorter sample starting in 2005, I show that the realized returns on cat bonds are not significantly correlated with the returns on equities, high-yield bonds, or mortgage backed securities.

The second assumption is that the excess return on cat bond market portfolio is correlated with the marginal rate of substitution of the specialist fund managers. This is likely to be the case due to the way the market is structured: majority of the securities are being held by funds that

[2]A frictionless macro-finance model predicts that the observations in Figure 1.1 line up on the x-axis.

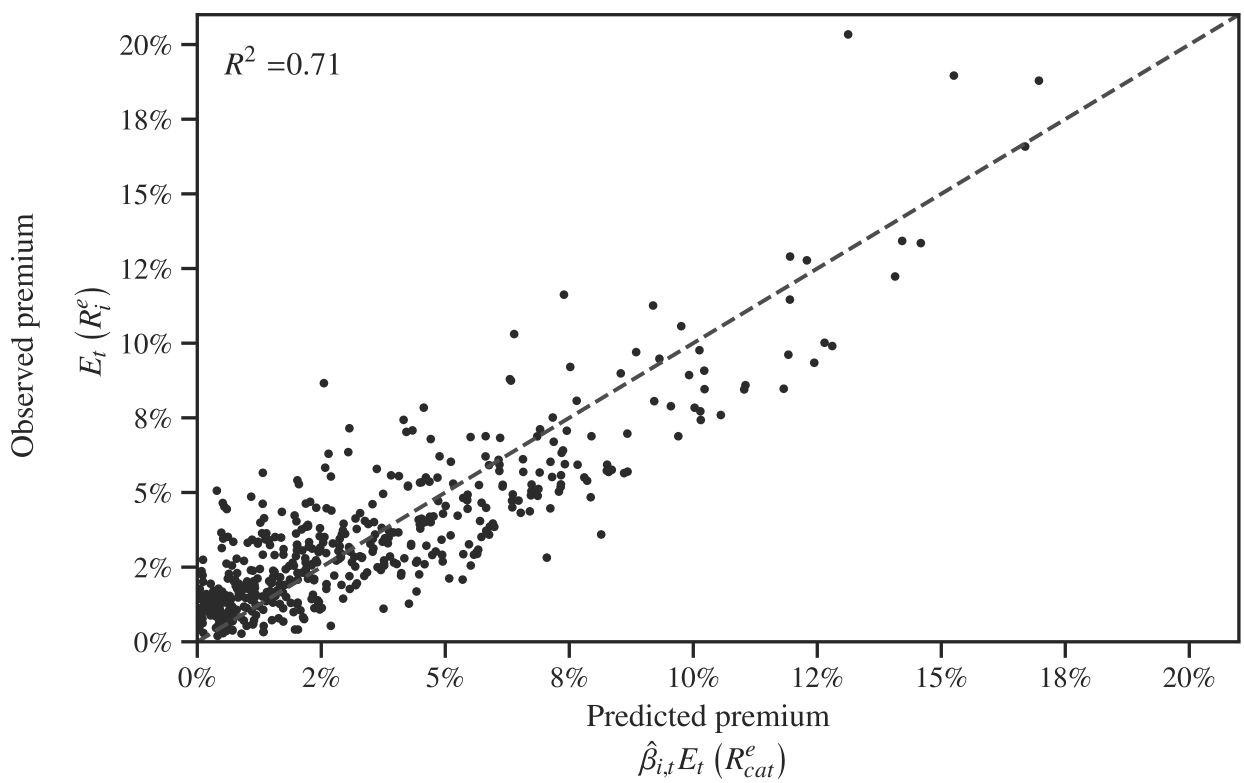

Figure 1.1: Observed vs. predicted risk premium.

Description: The figure plots expected excess returns (discount margin - expected actuarial loss) on sample cat bonds between 2003 and 2018 against their predicted values from a single-factor intermediary model. Bonds' betas are estimated from 500,000 years of simulated disaster data, and each dot represents a single cat bond observation in the end of June of a given year.

only specialize in this asset class[3], so their aggregate holdings must almost mechanically resemble a diversified asset class-specific portfolio. Indeed, I show that the correlation between quarterly returns on the specialist fund manager index, and the returns on my cat bond market portfolio is 0.87.

The third assumption is that the expected loss estimates from actuarial catastrophe risk models that I use to construct expected return measures are unbiased proxies for investors' expectations. If the estimates of these models were biased downwards, the differences between discount margins

[3] Aon Benfield (2018). The rest of the assets are being held by other institutions, such as mutual funds and hedge funds. Retail investors are virtually non-existent in the market due to regulatory constraints that mechanically limit their participation. For example, in the U.S. cat bonds are traded under the SEC 144A rule which implies that only qualified institutional buyers with more than \$100 million assets under management are allowed to invest in these securities.

and modeled expected losses could not be interpreted as expected excess returns. Instead, the observed spreads would be adjustments for the downward-biased loss estimates. I show that during my sample, realized losses are very close to the losses predicted by the actuarial models (albeit with large standard errors). Furthermore, if the observed risk premium on cat bonds was only due to downward-biased loss estimates, the probability that the cumulative losses on the bonds over the sample period would have been smaller or equal to its observed value is only 2.4%.

After having tested the cross-sectional implications of the intermediary model, I turn my attention to the time series properties of the aggregate premium on natural disaster risk and find several interesting patterns. First, during the early part of my 2003-2018 sample, the premium increases sharply after the occurrence of qualifying disasters, consistent with the findings of Froot and O'Connell (1999). Interestingly, however, I also find that the premium has decreased significantly after the financial crisis and seems to have become less responsive to the occurrence of disasters.

Second, consistent with the prediction of the intermediary model, this decrease in premium is associated with a gradual increase in available capital for the specialist funds relative to the size of the market. This increase in capital availability is consistent with the casual observation that after the financial crisis, there has been a gradual but large inflow of new institutional capital into the specialist funds, especially from pension funds that have tried to increase their yields in the low interest rate environment.

Third, a back-of-the-envelope calibration suggests that in order for the model to explain the aggregate premium on natural disaster risk during the sample, fund managers need to have a coefficient of relative risk aversion between 6.1 and 12.2. This suggests that intermediation frictions provide a quantitatively reasonable explanation for the observation that the compensation for natural disaster risk is too high compared to the level predicted by frictionless models.

Finally, I study alternative explanations for my cross-sectional and time series results. I find that the alternative hypotheses of Froot (2001): supply side market power, inefficient corporate form, liquidity and trading costs, and moral hazard and adverse selection are unlikely to explain the data. Furthermore, I show that the presence of some unobserved low probability disaster events that can cause large economy-wide disruptions (i.e., a peso problem) does not explain the pricing results.[4] Finally, I show that my result stays similar in a subset of bonds that are only exposed to the earthquake risk, suggesting that any potential exposure to climate change risk is not driving the result.

My paper contributes to several strands of literature. First, it is related to the empirical literature on intermediary asset pricing, in which several papers have shown that proxies for financial intermediaries' marginal rate of substitution have a high ability to price the cross-section of test assets across various asset classes such as equities and government bonds (Adrian, Etula, and Muir, 2014), options (Gârleanu, Pedersen, and Poteshman, 2009), and mortgage backed securities (Gabaix, Krishnamurthy, and Vigneron, 2007). He, Kelly, and Manela (2017) provides cross-sectional evidence from many asset classes and find strikingly similar prices of risk across the markets.[5] I complement this literature by testing this theoretical prediction in a laboratory where omitted risk factor concerns can be alleviated.

Next, following Koijen and Yogo (2015, 2016, 2018), there has been a resurgence of interest towards insurance markets among financial economists, mainly due to the sector's growing systemic importance and disruptions in business models due to financial innovation (e.g., Acharya, Philippon, and Richardson, 2016; Ge, 2019; Sen, 2019). However, the literature has mostly been focusing on life and health insurance markets, with the evidence being more limited in the property and casualty (P&C) markets. In this space, cat bonds represent a significant source of the ongoing disruption, and my results highlight the role of the institutions in understanding the pricing determinants of these emerging instruments.

Important previous work on the markets for catastrophe risk include Froot and O'Connell (1999, 2008) and Froot (2001) who show that the price of natural disaster risk in reinsurance markets is too high relative to a frictionless benchmark due to some supply side frictions. Fur-

[4]See Bekaert, Hodrick, and Marshall (2001) for discussion on a peso problem in the context of term structure anomalies, and Burnside et al. (2011) in the context of a carry trade.

[5]See also Chen, Joslin, and Ni (2019), Elenev, Landvoigt, and Van Nieuwerburgh (2018), Greenwood and Vayanos (2014), Haddad and Sraer (2019), Siriwardane (2019).

thermore, they develop various alternative explanations for this phenomenon. The main advantage of focusing on cat bond markets instead of traditional reinsurance markets is that the former are tradable securities with secondary markets instead of bilateral contracts between the buyer and the seller of protection. This makes it less likely that supply side market power is a key determinant of the observed premium.[6] Furthermore, the existence of secondary markets allows me to directly test and reject various alternative hypotheses. For example, I show that cross-sectional variation in cat bonds' liquidity (measured with secondary market trading data) is not a key determinant of the cross-sectional differences in expected returns. Instead, my results highlight that frictions in intermediaries' access to outside capital is key to understanding why sharing of natural disaster risk currently fails.[7]

Understanding the reasons for the failure to share natural disaster risk has also practical implications since climate change is likely to have starkly different impacts around the world. An emerging literature on climate finance studies the effects of climate change on financial markets, and how financial markets in turn can be utilized to mitigate the effects of climate change. For example, Bansal, Kiku, and Ochoa (2017) and Bolton and Kacperczyk (2019) study the implications of climate risk on asset prices.[8] One open question in this literature is how risks related to climate change can be most efficiently hedged and shared across agents. Understanding how natural disaster risks such as hurricane and wildfire risks can be efficiently shared is closely related to this agenda. Indeed, international organizations such as the IMF and the World Bank have emphasized the role of cat bonds in building resilience against natural disasters, especially in developing countries, in anticipation of the potential increase in quantity of such risks due to the climate change. My paper highlights the role of intermediaries as a key source of friction that currently inhibits risk sharing of catastrophic events.

The rest of the paper is organized as follows. Section 1.1 discusses the institutional background

[6] Indeed, Froot (2001) noted that once more data on cat bonds become available, this market provides a good setting to distinguish among different explanations.

[7] Other papers discussing alternative risk transfer mechanisms in catastrophe risk space Cummins, Lalonde, and Phillips (2004), Ibragimov, Jaffee, and Walden (2009), and Braun (2016).

[8] See also Daniel, Litterman, and Wagner (2018), and Hong, Li, and Xu (2019), Ilhan et al. (2020), and Krueger, Sautner, and Starks (2019).

of catastrophe bond markets and describes the data. Section 1.2 develops a simple intermediary asset pricing model and derives several pricing predictions. Section 1.3 develops empirical procedures to test these predictions. Section 1.4 provides the main results and evidence on the plausibility of the assumptions needed to rule out other risk-based stories. Section 1.5 discusses and tests various alternative hypotheses for the main results. Section 1.6 concludes.

1.1 Institutional background and data

In this section, I provide background information on how cat bonds are structured to establish that the bonds are plausibly not exposed to any traditional fixed-income risk. I also discuss how actuarial risk of these bonds is measured because it will be important in my empirical setting. Furthermore, I discuss investors in this market and show that the aggregate holdings of specialist funds seem to closely reassemble a market portfolio of the outstanding bonds. This observation motivates the theoretical setting in the next section. Finally, I discuss the construction and the key properties of the data.

1.1.1 Catastrophe bond markets

Background

The catastrophe bond market first started to develop in the late 1990s largely as a response to chronic lack of capital in the traditional reinsurance markets, especially after large natural disasters such as Hurricane Andrew in 1992 and the Northridge Earthquake in 1994. After these events, the shortage of risk capital caused reinsurance prices to rise significantly, which encouraged market participants to develop instruments to buy protection against natural disasters directly from the capital markets. Figure 1.8 shows the size of the U.S. dollar denominated public catastrophe bond market since 1997.

The largest peril category in this market has been tropical and subtropical cyclones, including hurricanes in North America, typhoons in Japan, tropical cyclones in Australia, and windstorms in Europe. Out of these, the largest subcategory has been U.S. hurricanes (especially in Florida and on the Gulf Coast), followed by windstorms in Europe. The second biggest peril category is earthquakes, focusing on regions such as the U.S. West Coast (especially California), the New Madrid Fault Line in the U.S. Midwest, Japan, and Mexico. Other typical perils include thunderstorms, winter storms and wildfires, and more exotic perils include volcanic eruptions and meteorite impacts. Figure 1.9 plots the ten most covered perils and geographies in my sample.

In addition to natural disasters, a similar securitization structure has been used for risks related to life insurance markets, such as longevity and mortality risks, and risk for pandemic outbursts. Furthermore, the structure has been used for some business risks such as motor insurance liability, offshore oil spill liability, liability due to large lotto jackpot winnings, and even the risk of forced cancellation of the 2006 World Cup in Germany. However, compared to natural disasters, these risk categories have been relatively small and are not part of my analysis.

Primary market structure

Figure 1.2 illustrates the typical structure of a cat bond issuance. A typical ceding party, or sponsor, is a (re)insurance company or a government entity who seeks protection against natural disasters from capital markets. It sets up a Special Purpose Vehicle (SPV) and enters into a reinsurance contract with it that transfers some natural disaster risks to the SPV in exchange for premiums. To be able to meet its potential obligations, the SPV issues one or multiple tranches of catastrophe bonds to investors so that its obligations and protection are fully matched. The SPV then invests the proceeds from the cat bond issuance in safe, liquid assets such as government bills through a collateral account. If no qualifying disasters occur, the sponsor pays frequent insurance premiums to the SPV, that in turn pays floating rate coupons to the bondholders. At maturity, the SPV liquidates the collateral account and returns the proceeds to the investors. However, if a qualifying event occurs, the SPV liquidates the account prematurely and pays the sponsor compensation for the damages, exhausting the capital in the SPV and causing the cat bonds to trigger.

A key advantage of the deal structure is that it eliminates the credit risk faced by the investors. If the sponsoring company gets into financial distress due to reasons unrelated to natural disasters, the principal in the collateral account stays intact and is returned to the investors at maturity, provided that no qualifying disasters occur. Similarly, because the bonds typically pay floating rate coupons, investors have little exposure to interest rate risk. More generally, one of the primary objectives of the deal structure is to minimize any traditional risks associated with fixed income securities to isolate the "pure play" natural disaster risk and offer investors assets whose payouts are uncorrelated with the rest of the economy.

Due to a favorable regulatory environment, these SPVs are typically domiciled in Bermuda or the Cayman Islands, and the cat bonds are listed in the stock exchange of the respective offshore location. However, the secondary market trading often takes place in the U.S. under the Rule 144A of the SEC. Oftentimes, one cat bond issuance consists of multiple securities or tranches with different levels of risk. For example, the issuance might consist of three $100 million tranches, where the first one covers losses from $200M to $300M, the next one from $300M to $400M, and the last one from $400M to $500M. It can also be that an issuance consists of multiple bonds that cover different perils or have different maturities. Figure 1.10 plots the number of issuances in my full sample with different number of tranches.

Risk modeling

Before the issuance, the sponsor generally hires a catastrophe risk modeling company to assess bond-specific actuarial risks. These commercial catastrophe models first emerged in the late 1980s to help insurance and reinsurance companies assess their exposures to large natural disasters. There are three risk modeling agencies that up to this point have been responsible for the risk-assessment of virtually all cat bonds: AIR, EQECAT, and RMS. Figure 1.11 shows the market shares of these companies in my sample.

A catastrophe model consists of four modules. A hazard module contains a large stochastic set of catastrophe events, which approximates all the possible disaster scenarios and their probabilities. For example, for a hurricane model, the hazard module consists of a large set of potential hurricane events with different paths and intensities. These modules are based on scientific research in the relevant fields and are built by scientists including climate scientists, seismologists, and geophysicists.

The inventory module contains geocoded information on all the insured assets that are exposed to potential catastrophes. For example, it can contain the list of all the insured properties, with detailed attributes such as replacement value, building material and soil type that are relevant when predicting damages due to catastrophic events.

The engineering module is a mapping from the parameters of a specific event and the attributes of a specific inventory item to predicted damages. For example, the module might predict that a hurricane with specific wind speed in a given location causes 50% damage to a wooden building in that location. Finally, the financial module maps expected damages to insurance losses by taking into account insurance policy details such as deductibles.

By simulating a large number of events from the hazard module, the model yields probability distributions on the frequency of catastrophic events and the associated financial losses. Based on these estimates, the risk modeler can assess the distribution of losses for a given cat bond. This loss distribution is typically summarized by three variables: attachment probability (annual probability that the bond loses the first dollar), annualized expected loss (expressed as a fraction of the face value), and exhaustion probability (annual probability of a full loss). It turns out that mathematically it is always the case that attachment probability $\geq$ expected loss $\geq$ exhaustion probability.

Often, a cat bond is also assigned a credit rating by a credit rating agency. The expected loss assessment of the risk modeling company is a key input in the process of determining this rating. However, it has recently become increasingly common that the sponsors don't seek to obtain a credit rating for their cat bonds, because the output from the risk models is deemed to be sufficient for investors to make informed assessments about the riskiness of the bonds.

Trigger types

One of the key features of any cat bond is its trigger type, which determines the conditions under which the bond pays out to the issuer. There are four main trigger categories: parametric, modeled loss, industry loss index, and indemnity. The main trade-off among these triggers is the basis risk for the issuer, and the speed and objectivity when determining whether a qualified event has occurred.

A bond with parametric trigger attaches if some observable conditions are met, such as wind speed exceeding some limit in a specific weather station, or ground acceleration being measured above a certain threshold. In general, this trigger type is considered to be the most objective and transparent, and determining whether a specific event qualifies is relatively fast. On the other hand, a parametric trigger can expose the issuer to a significant basis risk because the trigger is not perfectly correlated with the issuer's own exposures for which the bond is designed to provide a hedge.

A modeled loss trigger tries to reduce the basis risk by adding an extra layer to the process. After event parameters are measured, they are given as inputs to a risk model that transforms these inputs to company- or industry-specific loss estimates. The bond is then considered triggered if these modeled losses exceed some threshold.

Industry loss trigger is based on estimates of industry-wide losses due to a qualifying event. A service provider such as Property Claim Services (PCS) collects information on insurance companies' losses and generates an aggregated loss index. This further reduces the basis risk relative to the earlier trigger types, but it introduces a delay before the cat bond losses are determined after a potential event. This is because it takes time for the insurance companies to process claims, auditors to verify losses, and the index service provider to aggregate the information. For detailed discussion on this trigger type, see Cummins, Lalonde, and Phillips (2004).

Finally, an indemnity trigger basically eliminates the basis risk because the payoff is determined directly by the sponsor's losses due to a qualifying event. Again, the downside of the indemnity trigger is that it typically takes time before the losses are fully determined. Furthermore, an indemnity trigger introduces a risk for moral hazard: since the ceding company has protection against insurance losses, it can be more lenient in its payout policy (e.g., to promote customer loyalty). For this reason, indemnity bonds are typically associated with the issuing company retaining some proportion of the ceded risk.

Investors

Catastrophe bonds are a key constituent of a broader market segment called *alternative reinsurance capital*. In June 2018, the size of this market was approximately $98 billion, representing 16% of the global reinsurance capital. In addition to catastrophe bonds, the other major constituent is *collateralized reinsurance*. These direct reinsurance contracts—that are used to pass disaster risk from the cedent to the capital market investors—are more customizable but less liquid structures than cat bonds. Unfortunately, due to the more opaque nature of this market, no comprehensive data are available on these contracts. Cat bonds accounted for the majority of alternative capital outstanding until 2011, after which collateralized reinsurance contracts have gained significant market share. In 2018, cat bonds account for around a third of the capital outstanding, with collateralized reinsurance contracts accounting for the great majority of the rest.[9]

A majority of assets in the alternative capital sector are being held by specialist funds that focus on this particular market: the fact that the AUM of these funds is $99 billion[10]—the size of the whole market—provides *prima facie* evidence on their prominence. The major client group of these funds is pension funds especially in North America and Europe.

Figure 1.3 plots the quarterly returns on an index of these specialist funds against the returns on the market portfolio of outstanding cat bonds in my sample. The two series are highly correlated (0.87), suggesting that the holdings of these funds closely reassemble a diversified asset-class-specific market portfolio (which would obtain mechanically through market clearing if these funds were indeed holding all the assets). While in my empirical study I focus only on cat bonds (for which there are comprehensive data publicly available) instead of the whole alternative capital sec-

[9] Aon Benfield (2018).
[10] As of October 2019. Source: artemis.bm Insurance Linked Securities Investment Managers & Funds Directory.

tor, Figure 1.3 suggests that the unobserved collateralized reinsurance contracts have very similar payout structures than the observed cat bonds making a market portfolio of outstanding cat bonds a good proxy for the total portfolio of the specialist funds. Given this evidence, a structure in which direct holdings in the asset class are restricted to a group of market specialists while other investors must participate through them seems a good, abstracted description of the market.

1.1.2 Data

One of the significant obstacles for studying catastrophe bonds is that there are no standard data sources that have readily available, representative information on them. Hence, my sample construction begins by identifying the universe of cat bond issuances. For this purpose, I use the Catastrophe Bond & Insurance-Linked Securities Deal Directory of artemis.bm. Artemis is the leading news, analysis, and data media service that focuses on alternative risk transfer markets. Since its launch in 1999, Artemis has maintained a database that contains mostly qualitative information on all public and most private cat bond transactions. In total, there are 582 unique bond issuances in the database in January 2019 when the sample was constructed. Note that the total number of bonds is larger than the number of issuances because many bonds contain multiple tranches. On the other hand, the Artemis database contains bonds that are exposed to other types of risks than natural disasters (e.g., mortality risk). These bonds are excluded from the sample. Similarly, so-called private cat bonds that were not widely marketed are omitted due to the lack of publicly available data.

Next, I match cat bond issuances from Artemis with the CUSIP master file maintained by CUSIP Global Services. Since this database contains the whole universe of issued CUSIP identifiers, this matching enables me to obtain identifiers for all the issuers that have at some point requested one. This matching is done manually using issuer name and the time of the issue.[11] All the matches are then manually checked. Finally, ISIN numbers are manually searched from

[11] CUSIP master file contains the date when the CUSIP number was issued. Since most cat bonds are issued through a transaction-specific SPV, the identifier is typically issued during the same or the previous month compared to the issuance. Sometimes, the same SPV is used for follow-up issuances.

Bloomberg for bonds that do not have a CUSIP identifier. Overall, this method allows me to find an identifier for all but 10 non-private bonds (all these missing bonds are early issuances before my main sample starts in 2003).

Next, I obtain data on bond characteristics from Bloomberg, S&P credit ratings from Capital IQ, and Moody's and Fitch credit ratings from the websites of Moody's and Fitch, respectively. Then, I obtain information on covered perils and modeling companies from Artemis, and information on attachment probabilities, expected losses, exhaustion probabilities, and secondary market sheet prices from the annual insurance linked securities market outlook publication of Lane Financial LLC. Since 2000, this annual review has contained an exhaustive list of new issuances associated with data on these deal characteristics.

Secondary market transactions are obtained from TRACE and are available since 2002.[12] After cleaning the data from known errors and duplicate entries using the method of Dick-Nielsen (2009, 2014), I calculate daily prices as volume-weighted averages of individual trades. Then, I use the last available price observation per quarter to calculate discount margins. In Appendix A, I discuss how realized return on the bonds is measured.

Table 1.1 contains the summary statistics for the final sample. In total, it consists of 675 publicly issued cat bonds with CUSIP or ISIN identifiers. The average size (par value) of a bond is $130.8 million. The average time to maturity is approximately three years, and the average quoted spread is 7.3%. The average annualized attachment probability, expected loss, and exhaustion probability in the sample is 3.2%, 2.3%, and 1.8%, respectively. Said differently, these probabilities imply that the average bond in the sample would experience a trigger event once every 31 years and full loss once every 56 years. Note that if the bonds were actuarially fairly priced, the average spread should be close to the average expected loss. Consistent with the earlier literature, this is clearly not the case, and there is a large premium on natural disaster risk in my sample.

[12]Even though FINRA started to collect trading data on 144A corporate debt (most publicly traded cat bonds fall into this category) for regulatory purposes in 2002, it was not initially available for outside researchers. In 2013, FINRA filed a request to SEC to change the rule on dissemination of transactions, so that it would be able to provide trade information on 144A bonds to the public. The request was approved and 144A transactions were made available starting from July 2014. In the same SEC filing, FINRA asked for a permission to establish a historic data set for 144A transactions that would retroactively contain all the trading information since 2002.

Panel B contains variables for secondary markets. Based on turnover data from TRACE, cat bonds are thinly traded: 38.4% of bond-quarter observations are not associated with any trading activity. Out of the bonds that are traded, a typical annualized turnover is around 35.8%. Roughly speaking, this implies that after the issuance, a typical bond with 3-year time to maturity turns over once before it matures.

Throughout the analysis, I use two alternative measures of bonds' secondary market prices: average broker-dealer price quotes from pricing sheets and actual observed trade prices from TRACE. Both measures have their virtues and drawbacks. As discussed earlier, only 61.6% of bond-quarter observations are associated with at least one price observation from TRACE. Furthermore, the price might be stale if the last trade occurred early in the quarter, after which some relevant information on the bond became available. On the other hand, sheet prices are observed at the end of each quarter, so at least in principle they should contain all the relevant information up to that point. Furthermore, they are observable for virtually the whole sample of bonds.

However, as shown by Warga and Welch (1993), one must be careful when using indicative prices when studying asset pricing implications in the bond markets. More specifically, they show that measuring the timing and the magnitude of bond price reactions to new information is highly sensitive to the used price measure. The sheet prices are only quotes and can differ from the prices at which investors are willing to transact. This can bias any asset pricing results in unknown ways. For these reasons, I report all the results using both price measures. All the results generally agree when using either of these two measures, suggesting that potential problems associated with either of them don't have significant impact on the results. Since both measures agree, I use the sheet prices as my primary price measure as it has significantly broader coverage.

1.2 Theoretical framework

The key implication of intermediary asset pricing models such as He and Krishnamurthy (2013) is that financial intermediaries rather than households are marginal investors in capital markets, and that the marginal rate of substitution of these intermediaries prices assets. In the case of catastrophe bonds, I argue that the relevant intermediary is a specialized cat bond fund manager due to the earlier observations that their AUM is almost exactly the size of the outstanding instruments, and because their quarterly returns closely reassemble the quarterly returns on a portfolio of outstanding cat bonds in my sample. To formalize this hypothesis and to get testable predictions on equilibrium prices under this setting, I develop a simple model with frictions in financial intermediation using a market structure similar to Gabaix, Krishnamurthy, and Vigneron (2007).

1.2.1 Assets

Consider a two-period economy with $t \in \{0,1\}$. There exists a risk-free asset with perfectly elastic supply that pays a rate of return r. There can also exist an arbitrary number of other outside asset classes. At $t = 0$, N cat bonds are issued at par (price-per unit $P_{i,0} = 1$), with θ_i being the exogenous total amount issued for bond i. A bond promises to pay a per-unit coupon of $C_i = r + s_i$ at maturity ($t = 1$), where s_i is the endogenous spread. The bond has a chance to trigger due to an occurrence of a qualifying natural disaster, with the attachment probability being $\bar{x}_i$. The probability of full loss (exhaustion probability) is $\underline{x}_i$. Given $x_i \sim U(0,1)$, the value of the principal at maturity is

$$P_{i,1}(x_i) = \begin{cases} 1 & x_i > \bar{x}_i \\ F_i(x_i) & \underline{x}_i \leq x_i \leq \bar{x}_i \\ 0 & x_i < \underline{x}_i \end{cases}, \tag{1.1}$$

where $F_i\left(\underline{x}_i\right) = 0$, $F_i\left(\bar{x}_i\right) = 1$, and $F_i'(x_i) \geq 0$.

Given expected loss $el_i \equiv 1 - E_0\left(P_{i,1}(x_i)\right)$, the payout $X_{i,1}$ then simply is

$$X_{i,1} = P_{i,1}(x_i) + C_i\,, \tag{1.2}$$

and the expected excess return is

$$E_0\left(R_i^e\right) = s_i - el_i\,. \tag{1.3}$$

Note that if the bond was actuarially fairly priced, its quoted spread would be equal to the expected loss, and its expected return would be equal to the risk-free rate.

1.2.2 Investors

Frictionless benchmark

Before introducing intermediation frictions, it is interesting to consider how the assets would be priced in a frictionless benchmark case in which a representative household holds a well-diversified portfolio of all the available securities in the economy. In this case, the price of risk is the same across the markets, and a strictly positive stochastic discount factor (SDF) M_1 prices all the assets. The SDF is determined by the marginal rate of substitution of the representative agent, and it is a typically postulated to be a function of aggregate consumption, wealth, or other macroeconomically relevant state variables.

Natural disasters, on the other hand, are exogenous events whose occurrence is not affected by any economic shocks—at least in the short-term. Hence, in order for natural disasters to be correlated with any macroeconomic variables, there would need to be a direct causal link *from* the occurrence of natural disasters *to* the state of the aggregate economy. Clearly, a link between disasters and the real economy exists at a local level: a natural disaster can cause direct economic damages up to several hundreds of billions of dollars. However, on an aggregate scale, such a first-order causal link is unlikely to exist. First, these potential losses are several orders of magnitudes smaller than the value of global financial assets that are quoted in hundreds of trillions of dollars. Second, casual examination of the recent history does not reveal any clear causal link. For example, the most devastating hurricanes in the recent U.S. history occurred during economic booms (Katrina in 2005, Harvey, Irma and Maria in 2017) with no clear effect on the performance of the U.S. economy. Third, more systematic examination during the past century does not reveal any significant impulse responses of macroeconomic aggregates to the occurrence of natural disasters (See Section 1.4.3).

Given these considerations, I make the following assumption:

Assumption 1 (Marginal rate of substitution of a representative household is independent of the occurrence of natural disasters):

$$M_1 \perp x_i \; \forall i \,. \tag{1.4}$$

While this assumption is likely to hold for the types of events we have observed in the historical data, there is always a possibility that some rare events that are not observed in the historical data are large enough to actually cause economy-wide disruptions and hence to violate Assumption 1. Thus, I will relax this assumption further in Section 1.5.2 to rule out such alternative explanation.

Given Assumption 1, $E_0\left(M_1 R_i^e\right) = 0$ directly implies that

$$E_0\left(R_i^e\right) = 0 \,. \tag{1.5}$$

In other words, expected return on any cat bond is equal to the risk-free rate, and the spread s_i is equal to the expected loss el_i. Each cat bond is priced at its actuarially fair value because natural disaster risk is diversifiable in equilibrium.

As discussed earlier, the prediction in Eq. (1.5) is overwhelmingly rejected in the data. This can be due to three reasons. First, it is possible that investors true loss expectations $\hat{el}_i$ are higher than the ones obtained from the actuarial catastrophe models (el_i). If this was the case, an econometrician who uses el_i as a measure of expected losses would observe spread (s_i) being higher than el_i and would incorrectly reject the frictionless model. However, this is unlikely to be able to explain the cross-sectional evidence presented in this paper. Furthermore, in Section 1.4.4, I show that given the actual loss history, we can reject that the aggregate premium on cat bonds is only due to actuarial models underestimating the expected losses.

The second possibility is that our assumption that natural disasters are independent of the marginal utility of the representative household is wrong. While this is intuitively unlikely (we don't typically see big disasters being associated with economic recessions), I will return to this issue in Section 1.4.3 in which I show that in a panel of 13 developed countries and 66 years of macroeconomic and natural disaster data, I cannot reject the null hypothesis that natural disasters do not systematically affect consumption or GDP growth, house prices, or stock market returns.

The third possibility is that there are some market frictions that prevent the prices from adjusting to their efficient levels. In what follows, I pursue this explanation by introducing a friction that results in imperfect risk sharing.

Economy with financial intermediation

Assume that in addition to a representative household with high total wealth, whose marginal rate of substitution M_1 satisfies Assumption 1[13], the economy is also populated by F identical specialized fund managers with endowments of $w_{m,0}$. Outsiders cannot directly participate in the cat bond offerings; they can only invest in the asset class by giving funds to the market specialists.[14] To ensure that the managers invest prudently, the outsiders require the managers to have "skin in the game" by having their own wealth invested in the fund. More specifically, if w_0 is the total size of a fund at $t = 0$, at least $\alpha\%$ must come from the manager:

Assumption 2 (Fund managers have a limited capacity to raise outside equity):

$$w_{m,0} \geq \alpha w_0 \,. \tag{1.6}$$

Such capital constraint can be micro-founded, for example, by assuming moral hazard by the specialist (see He and Krishnamurthy, 2012).

The manager is risk averse and maximizes his utility over terminal wealth:

$$U(w_{m,1}) = E_0\left(w_{m,1}\right) - \frac{\rho}{2} Var\left(w_{m,1}\right) \,. \tag{1.7}$$

[13]Note that unlike Gabaix, Krishnamurthy, and Vigneron (2007), I don't assume that the outside investors are risk-neutral. As long as their marginal rate of substitution is independent of the occurrence of natural disasters, they can have arbitrary preferences.

[14]In the U.S., cat bonds are traded under the SEC Rule 144A that limit the participation to investors with more than \$100 million invested in securities. This rule automatically prevents for example retail investors from participating in these offerings.

The manager's problem is to choose the quantities q to build a portfolio of cat bonds that maximizes his expected utility.

1.2.3 Equilibrium

With an endowment of $w_{m,0}$, a manager is able to raise outside capital to have total assets under management of $\frac{w_{m,0}}{\alpha}$. Date 1 value of this portfolio is

$$w_1 = \sum_i q_i \left(P_{i,1}(x_i) + C_i\right) + \left(\frac{w_{m,0}}{\alpha} - \sum_i q_i\right)(1 + r)\,. \tag{1.8}$$

Similar to Gabaix, Krishnamurthy, and Vigneron (2007), I focus on a case where a fund manager is not capital-constrained and is able to purchase a desired portfolio of cat bonds, either because he is sufficiently wealthy or because α is sufficiently low. The first-order condition to the manager's optimization problem and market clearing ($Fq_i = \theta_i$, $\forall i$) results in a simple CAPM-like expression for the expected returns on cat bonds:

$$E_0(R_i^e) = \beta_i E_0(R_{cat}^e)\,, \tag{1.9}$$

where $E_0(R_{cat}^e)$ is the expected excess return on the cat bond market portfolio.[15] Even though natural disaster risk is diversifiable on aggregate, the fund managers who are marginal investors in this market end up holding an excessive amount of market-specific-risk which ends up being priced in equilibrium. Finding evidence that R_{cat}^e is a priced factor in the cross-section of catastrophe bonds would provide evidence for the hypothesis that the fund managers are marginal investors, because on the economy level, R_{cat}^e only consists of diversifiable risk that should not be priced.

The aggregate risk-premium is given by the following expression:

$$E_0(R_{cat}^e) = \alpha a \rho Var_0(R_{cat}^e)\,, \tag{1.10}$$

[15] See Appendix B for proof.

where a is the average position of a fund in the cat bond markets. We can see that the premium is proportional to the intermediary constraint α. If $\alpha = 0$, the managers are not constrained in their ability to raise outside capital. In such a case, the risk is shared efficiently across all the agents in the economy, and the pricing reflects the marginal rate of substitution of the outside investors. Accordingly, the cat bond market specific risk carries no premium. If, on the other hand, the intermediaries are constrained and $\alpha > 0$, the cat bond market specific risk has a positive price.

1.3 Empirical implementation

In this section, I develop an empirical approach to test the cross-sectional prediction given by Eq. (1.9). This basically requires empirical measures for expected excess returns and for betas with respect to a cat bond market portfolio for all test assets. A standard approach to test Eq. (1.9) would be a two-pass regression: first, estimate the betas from asset-specific time series regressions and then risk premia from a cross-sectional regression. The problem with implementing this approach, however, is that the sample for realized cat bond returns is relatively short, and the returns are highly skewed making such an approach infeasible.

Instead, I exploit a unique property of this market: compared to most other asset classes, forward-looking return distributions of the bonds are observable to a high degree of accuracy due to the existence of actuarial models that predict the probabilities of qualifying natural disasters occurring. Assuming these models provide unbiased proxies for investors' expectations, we can measure expected returns directly. Furthermore, combining these predicted return distributions with some simplifying assumptions on the correlation structure of bonds' exposures to different perils, I can estimate their betas using Monte Carlo simulations. This approach has the benefit that all premium and beta estimates are conditional—a feature whose importance has been emphasized by Jagannathan and Wang (1996) and Lettau and Ludvigson (2001), for example. In the following subsections, I discuss the construction of these measures more in detail.

1.3.1 Expected excess returns

The expression for expected excess return is given by Eq. (1.3): it is simply the difference between the quoted spread (s_i) and the actuarial expected loss (el_i). However, this expression assumes that the bond trades at par (i.e., there are no expected capital gains or losses due to price appreciation or depreciation) and that the actuarial expected loss is a good proxy for investors' conditional loss expectation. While these conditions are likely to hold at the time the bond is issued, several adjustments must be made when measuring expected excess returns in secondary markets. First, I use the bond's discount margin ($dm_{i,t}$) in place of the spread to take into account expected capital gains or losses when the bond is not trading at par. Second, because I observe the modeled loss estimates only at the time of the bond issuance, I use them as a proxy for expected losses also after the issuance.

One convenient feature of a typical cat bond is that at the end of each risk period, the contractual features of the bond are reset so that the conditional loss probabilities are again consistent with the unconditional ones. As a result, the unconditional modeled loss distribution is an accurate description of the conditional distribution once per year after the reset. This reset process is typically carried out by the original risk modeling company, who is often called "reset agent".

However, there are several reasons why using modeled loss probabilities as conditional loss probabilities might not be accurate between the resets. First, some bonds are structured to provide the sponsor coverage against its aggregate losses during a risk period (typically a year). If the losses have been accumulating faster or slower than originally expected, the conditional probability of a loss is different than the unconditional one. Second, risk modeling firms periodically update their models, and this can have a significant impact on the modeled loss estimates. Sometimes, these updates have had a big enough impact that have caused credit rating agencies to change their ratings on several bonds. A third major reason for a probability update is that a (potentially) qualifying disaster event has already occurred, making the conditional expected payoff significantly different from the unconditional one. The main effect of using a noisy measure of expected losses in my asset pricing tests is that it introduces measurement error to my dependent variable (expected excess return), decreasing the power of the test. Hence, I take several steps to alleviate this problem. While these adjustments help increase the power of the tests, none of them are essential and all the main results go through also without them.

Filtering distressed bonds

To reduce the effect of measurement error, I exclude those bonds from the sample whose conditional loss expectation is likely to differ significantly from the unconditional expectation. First, I drop any bond-quarter observation that is flagged as "distressed" in the secondary market pricing sheets. A bond becomes distressed after a (potentially) qualifying disaster event has occurred, but before the final losses are announced. In total, 7.4% of my sample bonds become distressed at some point. Second, I drop any bond after an event occurred that ultimately lead to some losses. Third, I drop any bond after a change in credit rating or if the bond is currently on credit watch by any of the three credit rating companies. As discussed earlier, the key determinant of a cat bond credit rating at the time of the issuance is its expected loss. If we observe a change in credit rating, it is likely that some material information has been released that has changed the conditional loss distribution. In total, 50.7% of my sample bonds are rated by S&P, 21.6% by Moody's, and 6.5% by Fitch. Finally, 64.1% of bonds are rated by at least one of the three agencies.

Finally, if possible, I limit my tests to the end of June each year. Some cat bond prices are highly seasonal due to seasonal patterns in the underlying event probabilities. For example, the majority of North Atlantic hurricanes and East Asia typhoons occur during the third quarter. Similarly, the European windstorm season takes place in winter. Accordingly, the most active quarter for new bond issuances is Q2, implying that for a typical bond providing aggregate coverage, the risk period resets during the second quarter. Due to these reasons, we expect the unconditional expected loss to be the best proxy for the conditional one at the end of Q2. Additionally, some bonds risk exposure ends before the maturity, making annual risk probabilities not applicable during the last year. Hence, I drop bonds that mature before the corresponding calendar quarter in the following year.

Variable reset

After 2013, a new contractual feature called variable reset became popular in new bond issuances. This feature allows the issuer to change the riskiness of the bond within specific bounds at the time of the reset so that the bond keeps occupying a specific layer of their "reinsurance tower", even if the issuer's portfolio has changed. As a consequence of any adjustment, the spread of the bond will be adjusted upwards or downwards to reflect the change in the underlying risk. While I do not observe the updated probabilities, I do observe changes in bonds' spreads. Hence, to adjust for variable resets, I adjust the actuarial probabilities every time I observe a change in a spread. More specifically, I first set $\bar{x}_{i,new} = \bar{x}_i \frac{s_{i,new}}{s_i}$, and then $el_{i,new} = el_i + \left(\bar{x}_{i,new} - \bar{x}_i\right)$ and $\underline{x}_{i,new} = \max\left(0, \underline{x}_i + \left(\bar{x}_{i,new} - \bar{x}_i\right)\right)$.

Hedging currency risk

Since some bonds in my sample are denominated in currencies other than USD (9.0% in EUR, 0.8% in JPY), they are exposed to currency risk from the perspective of a U.S. investor, which is potentially a systematic source of risk. To alleviate this problem, I hedge the currency risk of the principal using currency forwards. My results are not materially affected if I leave the currency risk unhedged or if I drop these bonds altogether from the sample.

Measuring risk-free rate

One additional assumption in Eq. (1.3) is that the reference rate on all bonds is some uniform risk-free rate instead of various different (albeit highly correlated) benchmark rates. In practice, there are several different reference rates for the bonds with the most typical ones being the three-month USD LIBOR and the three-month U.S. T-bill rate. Furthermore, a recent literature argues that interest rates on assets like U.S. treasuries reflect in part their money-like features such as liquidity and collateral value (e.g., Krishnamurthy and Vissing-Jorgensen, 2012; Nagel, 2016b). Hence, even if catastrophe bonds were actuarially fairly priced, we would not expect them to pay rates that are equal to the treasuries, but instead rates similar to other safe assets that lack the money-like features.

A novel way of estimating risk-free interest rates that are unaffected by the convenience yield on money-like assets is provided by van Binsbergen, Diamond, and Grotteria (2019). Given earlier considerations, I use their 12-month Box rate (reflecting 12-month holding period) as a risk-free rate in my main specifications. While this choice helps to reduce the unexplained level in my bonds' premia by 65bb, my results are similar if I use the more traditional U.S. government bill rate as a measure of risk-free rate instead.

1.3.2 Constructing modeled betas

In the following subsections, I develop a proxy for $\beta_{i,t}$ that is needed to test the relation given by Eq. (1.9) in the cross-section of sample bonds. Because of a short sample and highly skewed returns, we cannot estimate bonds' betas from the historical data of realized returns. Instead, I will estimate betas using simulated data and information on bonds' modeled loss distributions and underlying perils.[16] For each bond i, I observe three probabilities that summarize the annualized loss distribution: attachment probability ($\bar{x}_i$), expected loss (el_i), and exhaustion probability ($\underline{x}_i$). Furthermore, I observe which peril category and region the bond is exposed to. For single-peril bonds (i.e., bonds that are exposed to only one peril-region category) these categories are Atlantic Hurricane (North and Middle America), North America Earthquake, Pacific Hurricane (Middle America), Middle America Earthquake, South America Earthquake, Europe Windstorm, Europe Earthquake (including Turkey), Asia Typhoon, and Asia Earthquake.

Bond-specific return distribution

Let $x_i \sim U(0,1)$. x_i is a random variable that represents the severity of a particular disaster, with smaller values of x_i indicating a more severe disaster. For example, $x_i = 0.01$ implies a disaster that is expected to occur once every hundred years. Given a realization of x_i, I measure

[16] See Froot and O'Connell (1999) for somewhat similar approach to estimate catastrophe losses using Monte Carlo simulation.

the simulated realized excess return on bond i as:

$$R^e_{i,t+1}(x_i) = (1 + ref_{i,t} + dm_{i,t})P_i(x_i) - r_t\,, \tag{1.11}$$

where $ref_{i,t}$ is the benchmark rate of bond i, $dm_{i,t}$ is the discount margin and r_t is the risk-free rate.

One minor complication in measuring the return distributions is that I don't observe the shape of the modeled payoff function $P_i(x_i)$ in the domain of partial losses. Fortunately, the exact shape has little impact on the results. The simplest approach to approximate the function in the region of partial losses would be to assume that it is linearly increasing in x_i, or $P_i(x_i) = \left(\frac{x_i - \underline{x}_i}{\bar{x}_i - \underline{x}_i}\right)$, $x \in [\underline{x}_i, \bar{x}_i]$. However, this approach would result in $P_i(x_i)$ to have a different expected loss than the true function, unless the expected loss of i lies half-way between its attachment and exhaustion probabilities. To correct for this, I allow P_i to be concave or convex depending on the location of el_i:

$$P_i(x_i) = \begin{cases} 1, & x_i > \bar{x}_i \\ \left(\frac{x_i - \underline{x}_i}{\bar{x}_i - \underline{x}_i}\right)^{\phi_i}, & \underline{x}_i \le x_i \le \bar{x}_i \\ 0, & x_i < \underline{x}_i \end{cases}, \tag{1.12}$$

where $\phi_i = \frac{\bar{x}_i - \underline{x}_i}{\bar{x}_i - el_i} - 1$. In this case $P_i(x_i)$ has always the same expected loss as the true function.

Figure 1.12 illustrates $P_i(x_i)$ for three hypothetical bonds that all have exhaustion probability of 0.01 and attachment probability of 0.03. Bond A has expected loss of 0.02, which implies that $\phi_A = 1$ and $P_A(x_A)$ is linear in x_A in the domain of partial losses. If expected loss is closer to exhaustion probability than attachment probability, $\phi_i < 1$ and $P_i(x_i)$ takes a concave shape. This is illustrated by Bond B that has expected loss of 0.015. Accordingly, Bond C has expected loss of 0.025, implying $\phi_C > 1$ and convex shape for $P_C(x_C)$.

The first row of Table 1.2 summarizes the distribution of ϕ_i. The cross-sectional average of ϕ_i is 0.84, implying that for a typical bond, losses are a slightly concave function of the probability of the underlying peril. 25% of bonds have ϕ_i larger than or equal to one, implying a convex shape for the function. The actual functional form of P has little impact on the results. For example, setting $\phi_i = 1$ for all bonds (and disregarding the information in el_i) has little impact on any results that follow.

Correlation structure of losses

While I observe the loss distributions of individual bonds to a relatively high level of accuracy, the information on the correlation structure across bonds is significantly more limited. Hence, I make the following simplifying assumptions. First, I assume that disasters across peril categories are independent of each other: for any two bonds i and j in risk-categories c and d, I assume that $x_{i,c} \perp x_{j,d}$, $c \neq d$. This assumption implies, for example, that the occurrences of Japanese and U.S. earthquakes are drawn from two independent distributions, which is likely to be reasonable.

Second, for simplicity and lack of more granular data, I assume that the losses for bonds within the same peril category are drawn from a single distribution ($x_{i,c} = x_{j,c}$). This assumption overstates the correlations among bonds within same category, especially if the category consists of very heterogeneous bonds. For example, in the U.S. hurricane category, not all bonds are exposed to the same set of events that can affect the region. Some bonds are more exposed e.g., to Florida, which is the single most concentrated source of natural disaster risk, whereas other bonds are more exposed e.g., to New York, which is a less-covered area. As a result of this assumption, bonds that are exposed to a relatively isolated set of events within a large category have their beta estimates being overstated. Because the average of betas has to be close to 1, some other beta estimates are underestimated. When these betas are used as explanatory variables in pricing regressions, the fact that they are imprecisely estimated due to these simplifying assumptions on correlation structure leads to an errors-in-variables problem, and downward-biased estimates for the prices of risk. This bias works against me because it makes it more difficult to reject the null hypothesis that the cat-bond-market-specific risk is not priced.

Cat bond market portfolio

In an ideal setting, I would observe the peril exposures of all insurance-linked securities that form the market portfolio in the alternative capital sector. However, there are several data limitations. First, in addition to cat bonds, there are other types of instruments held by the market specialists (see Section 1.1.1). These include instruments such as collateralized reinsurance contracts, industry loss warrants, and sidecars. These other markets are significantly more opaque than the cat bond market, greatly limiting data availability.

Second, a significant proportion (49.5%) of cat bonds are so-called multi-peril bonds that are simultaneously exposed to several peril categories and geographies. For these bonds, the observed modeled loss distribution is an aggregate one over all the different perils. Since I cannot observe the contribution of different peril categories and geographies to the overall expected losses, I drop multi-peril bonds from the sample. While this is less relevant from the perspective of having fewer available tests assets, it affects the composition of the cat bond market portfolio.

If the distribution of peril exposures among collateralized reinsurance contracts and multi-peril bonds were similar to the composition among single-peril bonds, the betas estimated with respect to the single-peril cat bond portfolio would be similar to the true asset-class-specific betas. However, if the compositions differ, betas would be biased upwards for bonds whose peril categories are overrepresented in the single-peril bond portfolio, and downwards biased for underrepresented categories. This problem would be a manifestation of the traditional Roll (1977) critique, and it would likely to make it harder to reject the null hypothesis that the cat-bond-market-specific risk is not priced.

Simulation results

For each time period and peril category, I simulate 500,000 peril realizations, and calculate bond returns using Eq. (1.11). The return on cat bond market portfolio is a face-value weighted average of individual bonds' returns. Then, the bond's beta estimate is simply the regression slope of its excess returns on the excess returns on the cat bond market portfolio over the sample of

500,000 realizations.

Table 1.2 summarizes the results of these simulations, using either broker dealer sheet prices or observed prices from TRACE as a basis for the estimation. The reported figures come from a pooled sample of bonds' betas from 2003 to 2018 for β_{sheet}, and from 2005 to 2018 for β_{trace}. Both measures have very similar distributions, but β_{sheet} has 70% more observations due to better coverage of sheet prices compared to the actual ones.

The simulated betas range from 0.01 to 3.28. In general, a bond can have low beta for two reasons. First, if it provides coverage against a very remote risk layer with a low attachment probability, that is it is not exposed to events with relatively high probability of occurrence. As a result, it will have lower market beta compared to a bond that is also exposed to more frequent disasters from the same distribution.

Table 1.10 illustrates the two main sources of variation in simulated betas. It provides information of five selected bonds from two different issuances. Four of these bonds are in North Atlantic hurricane category (which is generally the largest one) whereas one bond is in Pacific hurricane category. The first source of variation in betas comes from differences in modeled losses. The first three bonds are all exposed to the same peril category, but because the third bond is significantly riskier than the two other bonds (measured with expected loss), it has a higher beta estimate. The first two bonds are quite similar in terms of modeled risk, and consequently their betas are also similar.

The second source of variation comes from bonds having exposures to different underlying perils. A bond that is exposed to a peril category with a larger amount of total risk-capital outstanding ends up having a higher beta (because the peril is more "systematic") compared to another bond that provides diversification benefits for the market portfolio. For example, the fourth bond in Table 1.10 is the riskiest on a standalone basis, but it has the lowest beta because Pacific hurricane is not a widely covered peril category in the catastrophe bond markets.

The fifth bond illustrates a source of measurement error in the betas. In my sample, this bond is allocated to North Atlantic hurricane category, but closer inspection reveals that its exposure is limited to Atlantic coast of Mexico. Because most of the exposure in this category comes from the U.S. East and Gulf coasts (especially from Florida)—as exemplified by the first three bonds—the fifth bond's beta is likely to be overestimated because it is calculated under the assumption that if the U.S. is hit by a hurricane, so is Mexico. Since I don't have a systematic way to measure exposures more granularly, some of the betas will be overestimated and others underestimated. Two last columns of Table 1.10 are related to pricing results that will be discussed in the next section.

1.4 Empirical evidence on the failure of risk sharing

In this section, I provide my main cross-sectional results that test Eq. (1.9). Furthermore, I discuss aggregate catastrophe risk and its time series properties in the light of Eq. (1.10). Finally, I provide validating evidence for the two main assumptions that are needed to interpret my main results as evidence for financial intermediaries being marginal investors in cat bond markets. First I show evidence that the occurrence of natural disasters has not systematically affected the macroeconomic state of a given country. Then, I show that the realized losses on my cat bond market portfolio during the sample period are very close to the expected losses given by the actuarial models providing validation for my practice of using actuarial models to construct expected return measure.

1.4.1 Main cross-sectional results

Table 1.3 reports the results of cross-sectional regressions of the following form:

$$E_t(R^e_{i,t+1}) = \lambda_{0,t} + \lambda_{cat,t}\hat{\beta}_{i,t} + \varepsilon_{i,t}\,, \tag{1.13}$$

where the dependent variable is the expected excess return on cat bond i in the end of June of year t (2003-2018), and $\hat{\beta}_{i,t}$ is the simulated beta estimate. All standard errors are clustered by bond issue.

In all 16 years, the estimate for $\lambda_{cat,t}$ is positive ranging between 1.02% and 7.62% implying that there has been a positive premium associated cat-bond-market-specific risk. All the t-statistics are highly positive, but they should be interpreted with caution in the early years due to relative small effective sample sizes. The last row reports the quarterly time series average of $\lambda_{cat,t}$ (2.06%) with the associated Fama and MacBeth (1973) t-statistic. Taken together, the results provide strong evidence against the null hypothesis that the asset-class-specific systematic risk is not priced. This evidence is inconsistent with a frictionless model in which only economy-wide aggregate risks is priced, but it is consistent with the alternative intermediary asset pricing model that allows market-specific risks to be priced if a market specialist is the marginal investor in the market.

While the previous test simply asks whether the proposed factor is priced, the next column labeled $\lambda_{cat,t} - E_t(R^e_{cat,t+1})$ reports the results of a more ambitious null hypothesis: estimated risk-premium is equal to its theoretically motivated value—the expected excess return on the cat bond market portfolio $E_t(R^e_{cat,t+1})$. While the main goal of this paper is to show that market-specific risks are priced, it is interesting to see how the simple model holds against this more stringent standard.[17]

Compared to the earlier test, the results are somewhat more mixed: in 6 out of 17 years, we cannot reject the null, whereas the time series test suggests a fairly strong rejection. Taken together, however, the results still generally favor the rejection, implying that the premium on market risk is not as high as expected based on theory. Note, however, that there is likely to be room for improvement for this particular model. As discussed throughout Section 1.3.2, the simulated beta estimates are likely to contain noise due to simplifying assumptions on the correlation structure of bonds' exposures and the fact that a significant portion of the market portfolio is unobserved. These errors are likely to lead to a downward bias in $\hat{\lambda}_{cat,t}$, making $\lambda_{cat,t} = E_t(R^e_{cat,t+1})$ easier to reject. Finally, the table reports the cross-sectional R^2 values for each year. They range between 0.29 and 0.84, suggesting a reasonably high fit.[18]

[17] As emphasized by Lewellen, Nagel, and Shanken (2010), taking this restriction imposed by theory seriously is an important part of any cross-sectional asset pricing test.

[18] Table 1.11 repeats the analysis with actual trade prices from TRACE instead of the sheet prices. The results are consistent with those in Table 1.3, with $\lambda_{cat,t} = 0$ being rejected for most of the years. Because of lack of trading

Table 1.10 illustrates how the measurement error affects the results. As discussed in Section 1.3.2, the beta of the last bond of the sample is likely to be overestimated because it was calculated under the assumption that the occurrence of hurricane events in the U.S. are perfectly correlated with those in the Atlantic coast of Mexico. This results in a high beta estimate and consequently high predicted premium (7.0%), which is significantly higher than the observed premium (2.8%). Figure 1.13 repeats the analysis presented in Figure 1.1 and highlights these five example observations, showing clearly the outlier bond.

As discussed earlier in 1.3.2, the variation in simulated betas comes mainly from two sources: in the same peril category, a bond with higher modeled losses end up having a higher beta. Between categories, the ones that are larger in terms of risk-capital outstanding (and hence more "systematic") have bonds with higher betas. Hence, it is interesting to see how well the model performs on these two different dimensions. Figure 1.13 provides these statistics. In particular, In addition to the overall R^2 (71%), I calculate "within category" R^2 that provides information on how much variation in expected returns the model explains within a particular risk category and year. "Between categories, within year" R^2 reports how much variation the model explains between categories due to their differences in average betas. Both of these figures are very similar at 0.63 and 0.62, suggesting that the model is successful at explaining variation in two distinct dimensions. Finally, "between years" R^2 is close to 1 because excess return on value-weighted market portfolio mechanically explains almost all time series variation in average excess returns.

1.4.2 Time series results

This section provides evidence on the time series evolution of the aggregate premium of natural disaster risk in a "back-of-the-envelope" context. In order to take the aggregate predictions of the intermediary model to the data, I follow Gabaix, Krishnamurthy, and Vigneron (2007) and translate the preferences of a mean-variance investor to the preferences of CRRA agent whose preferences

data, however, the sample only starts in 2005 and contains fewer observations in each cross-section. For this reason, the cross-sectional t-statistics are reported mainly for the sake of completeness. The fact that the results using actual prices are in strong agreement with the sheet price results suggest that the original results are likely not to be due to issues related to using bond prices that are not actual ones.

are

$$U(w_m) = \frac{w_m^{1-\hat{\rho}} - 1}{1 - \hat{\rho}}. \tag{1.14}$$

As shown by Gabaix, Krishnamurthy, and Vigneron (2007), such an investor has locally mean-variance preferences with risk tolerance of $\hat{\rho}/w_m$. Substituting Eq. (1.10) for this expression and rearranging gives:

$$\frac{E_t(R^e_{cat,t+1})}{Var_t(R^e_{cat,t+1})} = \hat{\rho}\frac{Size_t}{AUM_t}, \tag{1.15}$$

where $Size_t = \sum_i \theta_{i,t} P_{i,t}$ is the size of the cat bond market and $AUM_t = \sum_f w_{f,t}$ is the total assets under management of the specialist funds in year t. Intuitively, the more capital the fund managers have at their disposal relative to the size of the market, the lower the premium per unit of risk.

I measure the conditional variance in the left-hand side of Eq. (1.15) using the simulation framework introduced in Section 1.3.2. This gives me the observed premium per unit of risk. The right-hand side of Eq. (1.15) gives its predicted values. I measure $Size_t$ using estimates of Aon Benfield regarding the size of the alternative capital markets. The estimated AUM of the largest specialist funds is obtained from insurancelinked.com.

Next, I calibrate the model so that the average pricing error is zero by setting $\hat{\rho} = 6.1$, which gives us an estimate for the fund managers' coefficient of relative risk aversion. Note, however, that this estimate is obtained under the assumption that the specialist funds are the only investors in the alternative reinsurance capital markets which is counterfactual. While this assumption does not affect the time series variation of the predicted premium (assuming that their ownership share stays constant), it results in $\hat{\rho}$ being a lower-bound estimate for the managers' risk aversion. If, for example, the specialist funds held only 50% of the securities in this market, $\hat{\rho}$ would be equal to 12.2.[19]

[19]Unfortunately, the estimates on the specialist funds' market shares (and the definition these funds) vary across sources. For example, Aon Benfield (2009, 2018) report that the specialist catastrophe funds' share in the issuances they have been participating in is 40% in 2009 and 59% in 2018. On the other hand, Swiss Re (2012) reports that the share of dedicated funds is 56% and 61% in 2007 and 2012, respectively. Some other industry sources quote even higher figures: 70% (RMS, 2012) and 75% (Fermat Capital, 2015). To conclude, while it is difficult to obtain accurate information on the market shares of the specialist funds, the reported figures generally suggest that the share should be at least 50%.

Figure 1.4 plots the time series of observed and predicted premium per unit of risk. First, focusing on the observed premium, we can see that it is relatively speaking higher in the early sample and then gradually decreases after 2010. Furthermore, we can see that early in the sample, the premium seems to react to the occurrence of natural disasters: the premium jumps sharply upwards after 2005 and 2008 events. This is in line with Froot and O'Connell (1999) who find similar pattern in reinsurance markets between 1970 and 1994.[20] However, in the past few years, the reinsurance cycle in cat bond market seems to have weakened substantially. Most notably, the premium seems not to have reacted strongly to the record losses of 2017.

The predicted premium has a similar downward-sloping pattern as the observed premium. It increases after 2005 and 2008 events but not after the subsequent events, albeit the magnitude of the change is significantly smaller than for the observed premium. While limitations in data quantity and quality restricts the degree to which sharp conclusions can be drawn, the patterns are generally consistent with the prediction that the tightness of the specialist funds' equity constraint is a key determinant of the risk premium.

These patterns are not surprising to industry observers. After the financial crisis, there has been a gradual but large inflow of institutional capital into the specialist funds. It has been often said that reach-for-yield behavior in a low interest rate environment and investors gradually becoming more comfortable in investing in the alternative asset class as the market has matured and built some track record explain this shift. After 2017 losses, the funds were quickly able to raise new capital to replace losses, which contributed to the attenuated price reaction. Such a narrative is in line with the hypothesis of this paper that the market specialists are constrained in their ability to raise equity capital, and this has an effect on the equilibrium prices of the assets.

[20] See Duffie (2010) for discussion on how slow-moving capital can result in a sharp upward spike in reinsurance premium followed by a gradual reversal.

1.4.3 Empirical evidence on past natural disasters not having an aggregate impact on the economy

To motivate the identifying assumption (natural disasters are independent of marginal rate of substitution), I study the correlations between economic damages caused by natural disasters and selected macroeconomic variables. First, I obtain annual macroeconomic data on 13 developed countries[21] between 1950 and 2016 from Jordà-Schularick-Taylor Macrohistory Database (see Jordà, Schularick, and Taylor, 2017; Knoll, Schularick, and Steger, 2017). Furthermore, I obtain information on the economic damages caused by natural disasters in these countries from the Emergency Events Database (EM-DAT) of Centre for Research on the Epidemiology of Disasters (CRED).[22] The database contains event- and country-level information on mass disasters from 1900 to the present.[23]

I construct a measure of economic damages for each country by adding up all the economic damages caused by natural disasters in each year and country and scaling the measure by previous year's nominal GDP. More specifically, if we denote nominal dollar damages caused by disaster d in country i in year t with $DMG_{d,i,t}$, the annual measure of relative disaster damage in a given country is

$$dmg_{i,t} = \frac{\sum_d DMG_{d,i,t}}{GDP_{i,t-1}} . \tag{1.16}$$

To study the effect of natural disasters to macroeconomic variables, I estimate impulse-responses using Jordà (2005) local projections:

$$\Delta_h y_{i,t+h} = \gamma_i + \gamma_t + b_1 Small_{i,t} + b_2 Large_{i,t} + \varepsilon_{i,t} , \tag{1.17}$$

where $\Delta_h y_{i,t+h} = y_{i,t+h} - y_{i,t-1}$ is $h+1$ year change in the macroeconomic variable of interest,

[21] Australia, Belgium, Canada, France, Italy, Japan, Netherlands, Norway, Portugal, Spain, Switzerland, United Kingdom, United States.

[22] EM-DAT: The Emergency Events Database - Université catholique de Louvain (UCL) - CRED, D. Guha-Sapir - www.emdat.be, Brussels, Belgium.

[23] Since the data coverage prior to 1950 is relatively poor on majority of the panel countries, I focus on the post-1950 sample. The results are similar if we include all available disasters since 1900.

that include real consumption growth per capita, real GDP growth per capita, nominal house price index growth and return on stock market index. $Small_{i,t}$ is an indicator variable that is equal to 1 if $0.1\% \leq dmg_{i,t} \leq 1\%$. $Large_{i,t}$ is an indicator variable that is equal to 1 if $dmg_{i,t} > 1\%$. γ_i and γ_t are country and time fixed effects, respectively, that are included to increase the efficiency of the estimation but are not needed for causal interpretation, since the occurrence of natural disasters is arguably as good as random.

Figure 1.5 plots the results with shaded regions showing the 95% confidence interval on the estimates. In the case of small disasters, natural disaster damages have no significant effect on any variable of interest: the estimates are very close to zero with relatively tight confidence intervals (y-axes of the plots ares scaled to show ±2 standard deviation bounds for 1-year change in the y-variable of interest). For large disasters, the effects are still not significantly different from zero, with larger error bounds reflecting the relatively small number of qualifying disasters in the sample (17). Table 1.12 shows the regression results for $h = 0$ and $h = 2$. All R^2 are very small with or without fixed effects, and all coefficients are close to zero.[24] In total, these results provide support for the assumption that there is no first-order causal link from the occurrence of natural disasters to the macroeconomic performance.

1.4.4 Modeled vs. actual losses

Throughout the paper, I treat the estimated loss distributions from the actuarial catastrophe models as unbiased proxies for investors expectations. In this section, I provide evidence to support the plausibility of this assumption and rule out the alternative that all observed cat bond premium is due to modeled losses being downward biased. Note, however, that biasedness in itself would not be enough to explain the *cross-sectional* evidence presented in Figure 1.1 and Table 1.3. In order to explain the pattern, the estimates would need to be more downward biased for bonds with higher betas.

[24] Note that in the case of consumption growth, the 3-year effect is significant at 10% level but *positive*. Because it is hard to consider a reasonable channel why natural disasters should increase consumption, I interpret this result simply as evidence that natural disasters don't seem to have a significant negative effect on consumption.

To evaluate the accuracy of the actuarial models, I estimate the loss distributions on cat bonds using the simulation framework introduced in Section 1.3.2 for each year in the sample (2003-2018). Then, I calculate the cumulative loss distribution on a strategy that invest $1,000 on an equally weighted portfolio of single-peril cat bonds rebalanced annually at the end of June.[25] I assume that the occurrences of natural disasters are independent across years. While this assumption is likely to be reasonable regardless, it also simplifies the calculations significantly because the probability density function of sum of independent variables is simply a convolution of their density functions.

I obtain realized cat bond losses from artemis.bm. Note that for two large loss years 2017 and 2018, the final losses are not yet available for all bonds but instead estimated from observed market prices and other available information. Some bonds suffered losses both during 2017 and 2018. In these cases, it is not straightforward to attribute losses to distinct events, and I simply assume all the losses occurred already in 2017.

Panel A of Figure 1.6 plots the expected cumulative dollar losses implied by the actuarial models.[26] Actual losses are superimposed and shown in a bar graph. Three major cat events that contribute to the cumulative losses are the Tohoku Earthquake in Japan in 3/2011, the Chiapas Earthquake in Mexico and Hurricanes Harvey, Irma and Maria in 2017, and Hurricanes Florence and Michael in 2018.

Throughout the whole sample, the actual losses are well within the interdecile range shown in the plot. Note, that due to the earlier simplifying assumption that the exposures of all cat bonds within the same peril-geography category are perfectly correlated, these confidence intervals are overstated and provide an upper bound to the actual intervals. By the end of the sample in June 2019, the expected losses on the strategy were $198, compared to the slightly lower actual losses of $168 that is at the 43.6 percentile of expected loss distribution under the null that the actuarial models are unbiased. This result is highly consistent with the assumption that the actuarial

[25]Because we are interested in comparing the differences between predicted and observed losses, it is most appropriate to use equal weights both in cross-section and over time. This is because each bond, regardless of its size, should be equally informative about the validity of the actuarial models. Results are similar for a value-weighted portfolio.

[26]Figure 1.14 shows the results as a percentage of capital invested.

models are unbiased, although the confidence intervals are relatively large due to the highly skewed distributions of the events. For example, even if the assumption were true, we would need 289 years of additional data to reject at 10% level that the expected losses are biased more than 0.5% to either direction.[27] To sum up, while are clearly unable to reject the null that the modeled probabilities are unbiased, we also cannot reject that the estimated losses are biased by several percentage points due to the relatively large confidence intervals.

Next, I consider another null hypothesis discussed in Section 1.2.2: all observed risk premia on cat bonds are simply due to actuarial models being downward biased proxies of investors expectations that are based on true underlying probabilities. This is an important alternative hypothesis because if true, the anomalous pricing of natural disaster risk could be explained without introducing any capital market frictions.

To test this hypothesis, I shift the bonds' loss distributions upwards so that the expected losses are equal to the discount margins, as required by Eq. (1.5). Then, I repeat the earlier procedure and calculate cumulative losses on an equally weighted cat bond portfolio given these alternative probabilities.

Results are shown in Panel B of Figure 1.6. We can see that by 2008, the expected losses under this alternative became inconsistent with the observed losses and stayed that way afterwards. By the end of the sample in June 2019, there is only a 2.4% chance that the losses were smaller or equal to the observed ones if the null was true. Hence, we can reject the hypothesis that the observed premium on cat bonds is simply due to downward biased loss estimates. This result is to my knowledge new in the literature, and it provides support for the earlier papers who make the same assumption about premium not being due to downward biased loss estimates (e.g., Froot, 2001).

[27] Assuming market conditions stay similar to 2018.

1.5 Robustness and alternative explanations

In the previous section, I showed that the data are consistent with two key predictions of the intermediary asset pricing theory. First, in cross-section, bonds that have higher betas with respect to an asset-class-specific-market-portfolio have higher premiums. Second, in time series, there is a positive aggregate premium on natural disaster risk even though the risk seems to be diversifiable at the macroeconomic level. In this section, I discuss several alternative explanations to these facts and provide further empirical evidence to test their implications.

1.5.1 Froot (2001) hypotheses

A natural place to start the exploration of the alternative hypotheses is to study the ones developed by Froot (2001). In the context of traditional reinsurance markets, he shows that the premium on natural disaster risk is too high compared to the level implied by a frictionless benchmark model and develops hypotheses on supply side frictions to explain this finding.[28] The first of his hypotheses—insufficient reinsurance capital—is analogous to the intermediary story, so in the following subsections I focus on discussing the four other alternatives.

Liquidity and transaction costs

Out of all the hypotheses discussed in this section, liquidity premium is perhaps the most likely candidate *a priori* to fit both the cross-sectional and aggregate evidence in the cat bond markets. In particular, a rich literature in corporate bond markets has documented that bonds (especially high-yield bonds) with poor liquidity are traded at discount relative to the more liquid bonds (e.g., Bao, Pan, and Wang, 2011; Dick-Nielsen, Feldhütter, and Lando, 2012). Since cat bonds are also thinly traded in the secondary markets, it is plausible that they are also subject to a liquidity premium. If, for example, bonds that have higher betas also had poorer liquidity, my main cross-sectional results could potentially be explained by liquidity premium instead of compensation for risk due

[28] He also discusses three demand side hypotheses but notes that they cannot explain the pricing result because they predict a decrease instead of an increase in the premium.

to failure of risk-sharing between market specialists and other agents in the economy.

One challenge in studying liquidity premium is that several popular measures ironically require relatively high-frequency price data to be observable. Friewald, Jankowitsch, and Subrahmanyam (2012) consider several alternative liquidity measures that can also be calculated for bonds that are relatively illiquid. I implement all their measures that are feasible given the data limitations. The first set of proxies are based on TRACE trading data: quarterly turnover, number of trades, and average trading interval.[29] Somewhat more crude measures, based on bond characteristics, include issuance size, bond age, and time to maturity.

Table 1.4 shows the correlation matrix of simulated beta estimates and all liquidity measures. Consistent with Friewald, Jankowitsch, and Subrahmanyam (2012), I find that all the liquidity measures have very low correlations with one another. However, all measures also have very low correlations with the beta estimates, suggesting that cross-sectional differences in liquidity do not explain the cross-sectional differences in expected excess returns associated with differences in betas. Note also that unlike in equities, there is no material correlation between beta and size.

Even if the differences in liquidity are unlikely to explain why differences in betas predict expected returns, it is interesting to see if proxies for liquidity have explanatory power on expected returns on their own. Table 1.5 shows the results of Fama-MacBeth regressions where simulated market beta and a liquidity proxy are both used to explain the expected returns. As suggested by the low correlations between betas and liquidity measures, adding a second variable to the pricing regressions does not have a material impact on the estimated premium on market risk. Several liquidity measures seem to significant predictors of expected returns, however,. For example, bonds with higher turnover or more frequent trading seem to have lower expected excess returns, although the improvements on R^2 due to the addition of the second explanatory variable are relatively minor. Taken together, cross-sectional differences in liquidity do not seem to explain my main result, and even though liquidity may have some explanatory power, it does not seem to be a strong determinant of the expected returns in the cross-section of cat bonds.

[29] Number of trading days since last trade, calculated for each bond and trading day, and averaged over all trading days for a given bond and calendar quarter.

Inefficient corporate form

One puzzling phenomenon among reinsurance companies is that even though their liabilities are arguably idiosyncratic in nature and their assets are historically being held in the form of short-term notes and bills, the stock prices of publicly traded reinsurance companies have positive equity market betas. Regardless of what the underlying cause for this strange form of risk-transformation is, it can result in the cost of capital for reinsurance companies being above the risk-free rate, which in turn increases reinsurance prices above their actuarially fair values.

To rule out that a similar puzzling phenomenon plays a key role in the cat bond markets, I run regressions where I explain monthly excess returns on cat bond manager index[30] with excess returns on several other asset classes or strategies: equities, high-yield bonds, mortgage-backed securities and carry trade factor of Lustig, Roussanov, and Verdelhan (2011). Table 1.6 provides the results. Resulting beta estimates are generally very close to zero and insignificant, with R^2 ranging from 0.005 (MBS) to 0.026 (High-yield bonds). This result is perhaps not that surprising because the ability to provide returns that are uncorrelated with other asset classes is an integral part of cat bond funds' value proposition. This result suggests that an end investor who can access cat bonds only through a specialist fund can reasonably expect returns that are uncorrelated with major asset classes in his portfolio, and hence the investor should be willing to allocate capital to such a fund without imposing high requirements on the expected rate of return.

Moral hazard and adverse selection

One potential explanation for the observed premiums is the possible presence of moral hazard or adverse selection. In particular, if a cat bond has an indemnity trigger,[31] the issuing company might have an incentive to increase the riskiness of its liabilities after the bond issuance. Similarly, it is possible that companies who have superior information (compared to the risk modeling companies) on the riskiness of their liabilities are more likely to issue bonds, so that on average

[30]Eurekahedge ILS Advisers Index.

[31]Bond triggers if the underwriting losses of the issuing (re)insurance company are above some threshold, see Section 1.1.1 for description of the trigger types.

the modeled risk of the bonds is biased downwards. In both of these cases, investors are likely to demand a higher premium for their holdings to adjust for these frictions.

Among practitioners, such concerns have been heavily discussed, and several measures in the securities design have been taken to mitigate these problems. First, the terms of the bonds are generally adjusted annually to take into account any changes in the riskiness of the underlying liabilities, so that the actuarial loss probabilities are at the same level as when the bonds were first issued. This "reset" is typically carried out by the same modeling company that was involved in the issuance. The second measure to mitigate moral hazard is that whenever an indemnity bond is issued, the issuing company typically retains some proportion of the ceded risk.

Despite these measures, it is still possible that there is room for such frictions to affect the pricing of the bonds. If, for example, these frictions were more pronounced among bonds with higher modeled risks and bonds that are exposed to larger risk categories, they could potentially explain my main cross-sectional findings. To study this issue, I repeat my main analysis in a subsample of bonds whose triggers are not a function of the actual losses of the issuer but that of an objective and observable physical conditions (e.g., wind speed exceeding some limit in a specific weather station). These bonds are not likely to be exposed to these frictions, because the issuers cannot influence the riskiness of these bonds after the issuance, and it is unlikely that the issuers have superior information on the actuarial probabilities of the underlying disasters occurring compared to a specialized risk modeling company.

The last column of Table 1.7 shows the Fama-MacBeth estimates for a subsample of bonds with parametric or modeled loss triggers. This restriction does not have a material impact on the results, implying a rejection of the hypothesis that that moral hazard or adverse selection is a key determinant of the pricing results.

Market power

There is ample anecdotal evidence in traditional reinsurance markets that the reinsurance companies exert market power and are able to keep premiums high in a setting that reassembles bilateral bargaining. In such cases, the premium on natural disaster risk can be above the marginal cost of providing the capital.

While such a hypothesis is likely to help explain the prices in traditional reinsurance markets, it is unclear how it would apply to cat bond markets. Typically, dozens of investors participate in any given bond issuance, limiting the scope for any single fund manager to limit participation in an attempt to influence the price. Even if it was the case that the specialist fund managers were able to exert some market power and obtain yields that are above their marginal cost in the primary markets, it is unclear how such a hypothesis can explain that the prices remain at a similar level and are relatively sticky also in the secondary markets. Also, it is unclear how such a hypothesis would explain the cross-sectional results.

1.5.2 Peso problem

In this section, I study whether we can explain cat bond prices with a frictionless framework that accounts for the possibility of rare, high-impact tail events. To do this, I relax Assumption 1 that natural disaster risks are diversifiable and not priced. Instead, I only assume that "small and frequent disasters" are diversifiable while allowing "extreme and unobserved disasters" to have arbitrary risk premium to reflect the possibility that there exists some major disaster states that we have not observed in the historical data. These unobserved events could potentially disrupt the whole economy to the extent that these states are associated with a large risk-premium in a frictionless rare disasters framework such as Barro (2006).

The details of the model are discussed in Appendix C, but its key prediction is that even if rare tail events are priced, a portfolio that buys a riskier tranche and shorts a safer tranche of the same bond issuance should have no risk-premium if the safer bond is still risky enough that it fully triggers before some "rare event threshold" is reached. Intuitively, by comparing the pricing of two tranches of the same bond issuance, we can difference out the effect of potential peso states, and focus on the pricing of the high-frequency disasters. Under the assumption that these small disasters are not priced, the difference in the bonds' spreads should be equal to the difference in their expected losses. Note that in this setting the term "small disasters" should be interpreted broadly: these events can cause large local damages, but are not big enough to cause economy-wide disruptions. Generally speaking, the natural disasters we have observed in the panel of 13 developed economies since 1950 are likely to fall into this category, because these events have not been associated with large drops in GDP or aggregate consumption.

Figure 1.7 plots the differences in spreads ($\Delta s_{i,j}$) and expected losses ($\Delta el_{i,j}$) between safe and risky tranches among my sample cat bond issuances.[32] The prediction of the null model is that the observations should line up to the 45-degree line, so that the differences in spreads would be equal to the differences in expected losses. Clearly, this is not the case. Instead, the differences in spreads seem to increase faster than the differences in expected losses. This suggests that the markets demand a premium for carrying an extra unit of non-extreme catastrophe risk.

To test this prediction more formally, I estimate the following model:

$$\Delta s_{i,j} = \lambda \Delta el_{i,j} + \varepsilon_{i,j}\,, \tag{1.18}$$

where each i denotes a distinct pair of consecutive tranches in a cat bond issuance j. Crucially for my analysis, all other characteristics, such as peril categories, maturities, and trigger types remain the same for the different tranches of the same issuance, implying that they are exposed to the same set of events but with different sensitivities. Note that if an issuance j consists of more than two tranches, one tranche can be simultaneously the safe bond in pair i and the risky bond in pair i'. Because pricing errors in these pairs are likely to be correlated, I cluster the standard errors at the issuance (j) level in the analysis that follows.[33] If our frictionless benchmark model holds, we should find that $\lambda = 1$.

Note that the null hypothesis ($\lambda = 1$) only holds for those pairs of cat bond tranches in which the safe tranche is risky enough that it is likely to fully trigger in any priced peso state. For example,

[32]Note that in this setting, we don't need detailed information on the correlation structure of different bonds because we are not estimating any betas. As a result, we can also include multi-peril bonds in the analysis which helps increase the sample size.

[33]The results are robust to clustering standard errors at month-level, although this specification amplifies the potential small-sample issues for most restricted specifications.

comparing a bond that triggers only if a 1-in-10,000 year event occurs to a bond that triggers due to a 1-in-1,000 year event is probably not reasonable because it is possible that an event that is big enough to be priced is more frequent than these extreme cases. On the other hand, if we compare a 1-in-100 year event to 1-in-50 year event, it is unlikely that these events are priced because historically we haven't see an indication that such disasters have had a significant macroeconomic impact.

Table 1.8 shows the formal results for an OLS estimation of Eq. (1.18). I start in the first column by including all the observations in the sample, regardless of the expected losses on the safer tranche (el_A). Then, I proceed by gradually restricting the sample more by omitting observations with small el_A. The estimates of λ seem to stay remarkably stable around 1.5 regardless of the minimum value of el_A in sample selection.[34] However, due to the small number of observations and clusters in columns (5) and (6), the standard errors in these columns should be interpreted with caution.

These estimates imply that for every percentage point increase in the probability of loss, the cat bond's spread increases by 1.5 percentage points, an economically significant premium over the actuarially fair price. These findings imply a rejection of a frictionless benchmark model in which frequent disasters are not priced, even if there is a possibility of rare peso states with high state prices.

1.5.3 Probability weighting

Barberis and Huang (2008) show that if market participants overweight extreme probabilities, assets with negative skewness (*not* coskewness) will have positive excess returns. Furthermore, if there are N such assets in the economy, one group of otherwise identical investors will buy all the assets forming a portfolio, while N groups will each short sell one such asset. These predictions are consistent with several features of the cat bond data. First, even if returns on cat bonds are not

[34] While cat bonds generally stay on the market until the maturity, some bonds include an option for the issuer to redeem them prematurely. Because this decision is potentially contingent on the state of the economy, I repeat the analysis by omitting callable bonds from the analysis and obtain similar results.

correlated with other assets in the economy, they will still have high returns due to their negative idiosyncratic skewness. Furthermore, a market structure where several issues each issue one (or a few) bonds and sell them to specialist funds that hold diversified portfolios of them seems to be consistent with the prediction.

In the cross-section of expected returns, however, the model's prediction goes to the wrong direction: less negatively skewed bonds have higher expected returns. Basically, a cat bond with higher expected loss has a less negatively skewed return distribution, and based on earlier results higher expected losses are associated with higher expected returns. Table 1.9 shows the results of Fama-MacBeth regressions where bonds' market betas and skewness (both obtained from simulation framework described in Section 1.3.2) are both used to explain the expected returns. As a stand-alone variable, return skewness has positive price, but this significance is mostly subsumed by market beta due to their high positive correlation. Adding skewness as a second expanatory variable adds little expanatory power compared to the baseline one-factor model with market betas.

1.5.4 Other explanations

Table 1.7 shows the Fama-MacBeth estimates for several subsamples of bonds. First, approximately 9% of bond-quarter observations are associated with bonds that are callable. Although this option is rarely exercised in practice (remember, the bonds typically pay floating rate coupons), its presence could have a meaningful impact on the prices of the associated bonds. Hence, I consider a sample where such callable bonds are excluded. Second, it is in principle possible that bonds associated with hurricane risk are exposed to climate change risk in unknown ways, and such risk could be systematic at macroeconomic level. Although I consider this unlikely because the term of the bonds rarely exceeds 3 years—which is likely to be too short time window for a potential climate change risk to materialize—I repeat the analysis only for earthquake bonds whose cash flows are not likely to be affected by the climate change. Finally, the last subsample test is associated with bonds that have a parametric and modeled loss trigger (discussed in Section 1.5.1). None of these sample restrictions have a material impact on the results: the average premium estimate is not materially affected and stays at around the 2% level.

1.6 Conclusions

I studied the pricing of catastrophe bonds and showed that the majority of the variation in their expected excess returns can be explained with a simple one-factor intermediary asset pricing model. Because natural disaster risks are arguably uncorrelated with macroeconomic fluctuations, this observation is inconsistent with frictionless benchmark models, but consistent with models of market segmentation and financial intermediation where asset-class-specific-risks can be priced even if they wash out in the aggregate. This market friction implies a failure of risk sharing.

What are the potential reasons that prevent capital from flowing to the catastrophe risk market and resolving the risk sharing problem? Based on discussions with market participants, one major barrier of entry is that the majority of institutional investors are unfamiliar with the structure and main properties of this market. Every time a specialist fund wants to attract a new investor, it must spend a considerable effort to educate the outside investor about the key properties of the market, such as risks and their modeling. For new players interested in entering the market, a major roadblock is the high fixed cost associated with obtaining required capabilities and granular data. Given the market's relatively small size, these fixed costs of money and effort cause many institutions to shy away from this market.

Having established that the price of catastrophe risk is too high due to imperfect risk sharing, the next natural question is: what would be the likely consequences of alleviating these frictions? There is a globally large insurance coverage gap—especially in the developing countries—that is another manifestation of the failure in risk sharing. If reinsurance prices were to converge closer to their actuarially fair values, how would it affect the cost of issuing policies in the primary insurance markets, and what would be the potential welfare gains to the end user of insurance protection? I plan to explore these questions in future research.

1.7 Tables and figures

Table 1.1: Summary statistics

Panel A presents the summary statistics for US dollar denominated catastrophe bond primary market between 1997 and 2018. Size denotes the par value at the time of the issuance. Attachment probability is the estimated annual probability that the cat bond loses its first dollar. Exhaustion probability is the estimated annual probability that the cat bond fully defaults.
Panel B presents the summary statistics for secondary markets for the same sample of bonds. Turnover is the quarterly sum of daily trading volumes (from TRACE) divided by par value, and is expressed in annualized basis. 38.4% of bond-quarter observations are associated with no trading. Discount margin$_{sheet}$ is the average price quote (expressed as discount margin) from several broker-dealers, observed in the end of each quarter since June 2002. Discount margin$_{trace}$ is calculated each quarter using the last daily volume-weighted average trade price observation from TRACE since December 2002. Largest 1% of observations are trimmed for both yield variables.

Panel A: Bond characteristics (primary market)

Variable	N	Mean	St Dev	Min	25%	50%	75%	Max
Size ($ million)	675	130.8	119.6	1.8	50.0	100.0	175.0	1500.0
Term (months)	675	36.9	13.1	5.0	35.4	36.3	47.6	120.0
Spread (%)	658	7.3	5.1	0.7	4.0	6.0	9.3	49.2
Attachment probability (%)	657	3.2	3.2	0.0	1.1	2.0	4.2	23.2
Expected loss (%)	661	2.3	2.3	0.0	0.9	1.5	3.1	15.8
Exhaustion probability (%)	656	1.8	1.8	0.0	0.6	1.1	2.3	12.0

Panel B: Secondary market

Variable	N	Mean	St Dev	Min	25%	50%	75%	Max
Turnover (%)	5,969	35.8	79.1	0.0	0.0	8.0	37.9	2000.0
Discount margin$_{sheet}$ (%)	6,538	6.6	4.8	0.0	3.4	5.3	8.1	39.5
Discount margin$_{trace}$ (%)	3,994	6.7	4.6	-0.3	3.5	5.4	8.3	36.9

Table 1.2: Summary of simulation results

The table presents the summary statistics for Monte Carlo simulations that are used to estimate betas. In the end of June between 2003 and 2018, betas are estimated by drawing 500,000 realizations from the loss distributions of each single-peril cat bond, assuming distributions are perfectly correlated among bonds within the same peril-geography category and not correlated between the categories. Available categories are Atlantic Hurricane (North and Middle America), North America Earthquake, Pacific Hurricane (Middle America), Middle America Earthquake, South America Earthquake, Europe Windstorm, Europe Earthquake (including Turkey), Asia Typhoon, and Asia Earthquake. ϕ is a parameter that controls the shape of the assumed loss distribution, with values larger (smaller) than one implying convex (concave) shape in the domain of partial losses. $\phi = 1$ implies loss function is linear in event probability. $\hat{\beta}_{sheet}$ is the beta estimate using sheet prices. $\hat{\beta}_{trace}$ is the beta estimate using actual trading prices from TRACE (data available between 2005 and 2018).

Variable	N	Mean	St Dev	Min	25%	50%	75%	Max
ϕ	2,075	0.84	0.27	0.05	0.69	0.82	1.00	3.00
$\hat{\beta}_{sheet}$	2,158	1.04	0.71	0.01	0.42	0.92	1.54	3.28
$\hat{\beta}_{trace}$	1,267	0.98	0.69	0.01	0.36	0.86	1.53	2.94
N_{trials}	500,000							
N_{perils}	9							

Table 1.3: Pricing of catastrophe market risk

The table presents estimation results for cross-sectional regressions of the form

$$E_t(R^e_{i,t+1}) = \lambda_{0,t} + \lambda_{cat,t}\hat{\beta}_{i,t} + \varepsilon_{i,t} \,,$$

where the dependent variable is the expected excess return on cat bond i in the end of June of year t, and $\hat{\beta}_{i,t}$ is the simulated beta estimate. Standard errors are clustered by bond issue. Column $\lambda_{cat,t} - E_t(R^e_{cat,t+1})$ reports the difference between the premium estimate and its theoretical value: excess return on cat bond market portfolio. The last row contains quarterly time series averages of different parameter estimates, with associated Fama and MacBeth (1973) standard errors. R^2 in the last row is the time series average of cross-sectional figures. All prices are calculated using sheet prices. Bolded coefficients are significant at 1% level.

t	$\lambda_{0,t}$	(t-stat)	$\lambda_{cat,t}$	(t-stat)	$\lambda_{cat,t} - E_t\left(R^e_{cat,t+1}\right)$	(t-stat)	R^2	N	$N_{clusters}$
2003	**1.47**	16.98	**2.14**	17.42	**-1.45**	-11.78	0.73	30	12
2004	0.09	0.12	**1.54**	3.11	-0.31	-0.63	0.51	36	18
2005	**0.84**	6.56	**1.09**	12.08	**-0.88**	-9.72	0.42	34	16
2006	-2.51	-2.54	**7.62**	9.62	2.13	2.69	0.82	33	18
2007	**1.49**	3.10	**3.78**	5.01	-0.96	-1.27	0.71	40	28
2008	**1.53**	4.88	**2.86**	8.11	**-1.18**	-3.35	0.72	33	27
2009	**3.29**	5.10	**4.03**	5.14	**-2.97**	-3.79	0.71	22	17
2010	**3.10**	5.51	**1.99**	5.50	**-2.86**	-7.90	0.53	30	21
2011	1.07	1.25	2.62	2.54	-0.77	-0.75	0.42	22	15
2012	**1.21**	3.23	**4.08**	11.69	**-1.58**	-4.54	0.84	31	27
2013	**0.79**	3.75	**2.17**	8.52	**-1.02**	-4.02	0.76	42	35
2014	**1.15**	6.20	**1.39**	5.09	**-1.22**	-4.45	0.54	48	39
2015	**1.09**	7.04	**1.23**	6.85	**-1.12**	-6.22	0.60	50	39
2016	**0.90**	5.56	**1.02**	5.28	**-0.70**	-3.65	0.53	40	29
2017	0.53	2.38	**1.21**	3.64	-0.08	-0.25	0.31	46	32
2018	0.35	1.21	1.15	2.56	0.08	0.17	0.29	44	31
FM	**1.23**	9.41	**2.06**	11.67	**-1.10**	-9.02	0.49	63	

Table 1.4: Correlation of beta and liquidity measures

The table presents correlations of simulated betas and various liquidity measures of Friewald, Jankowitsch, and Subrahmanyam (2012). Annualized turnover, number quarterly trades, and trading interval (number of days since the bond last traded) are measured from TRACE. Characteristics-based measures include size (amount issued), age, and time to maturity. All liquidity measures are scaled by their pooled standard deviation.

	$\hat{\beta}$	Turnover	N trades	Trd intvl	Size	Age	Maturity
$\hat{\beta}$	1.00	0.06	0.03	-0.04	-0.01	-0.15	-0.14
Turnover	0.06	1.00	0.65	-0.10	0.04	-0.10	0.13
N trades	0.03	0.65	1.00	-0.25	0.44	-0.11	0.15
Trd intvl	-0.04	-0.10	-0.25	1.00	-0.21	0.31	-0.18
Size	-0.01	0.04	0.44	-0.21	1.00	-0.04	0.12
Age	-0.15	-0.10	-0.11	0.31	-0.04	1.00	-0.57
Maturity	-0.14	0.13	0.15	-0.18	0.12	-0.57	1.00

Table 1.5: Pricing of catastrophe market risk and liquidity

The table presents estimation results for quarterly Fama-MacBeth regressions of the form

$$E_t(R^e_{i,t+1}) = \lambda_{0,t} + \lambda_{cat,t}\hat{\beta}_{i,t} + \lambda_{liq,t}LIQ_{i,t} + \varepsilon_{i,t}\,,$$

where the dependent variable is the expected excess return on cat bond i in the end of quarter t. $\hat{\beta}_{i,t}$ is the simulated beta estimate, and $LIQ_{i,t}$ is the liquidity proxy in the given regression. The measures are annualized turnover, number quarterly trades, and trading interval (number of days since the bond last traded), size (amount issued), age, and time to maturity. All liquidity measures are scaled by their pooled standard deviation. R^2 is the time series average of cross-sectional figures. All prices are calculated using sheet prices. *, **, and *** denote statistical significance at 10%, 5%, and 1% level, respectively.

	Liquidity measure (LIQ)					
	Turnover	N trades	Trd intvl	Size	Age	Maturity
λ_0	1.34***	1.24***	1.36***	1.29***	1.40***	1.39***
	(0.15)	(0.15)	(0.14)	(0.14)	(0.15)	(0.15)
λ_{cat}	2.03***	2.06***	1.81***	2.08***	1.99***	2.00***
	(0.20)	(0.20)	(0.17)	(0.20)	(0.20)	(0.20)
λ_{liq}	−0.20**	−0.37***	−0.10	−0.30***	−0.15***	0.03
	(0.09)	(0.10)	(0.10)	(0.08)	(0.05)	(0.05)
N	57	57	57	57	57	57
R^2	0.54	0.54	0.53	0.55	0.54	0.54

Table 1.6: Explaining realized cat bond manager returns with other asset classes

The table presents estimation results for time series regressions of the form

$$R^e_{ILS,t} = a_i + b_i R^e_{i,t} + \varepsilon_{i,t}\,,$$

where the dependent variable is monthly excess return on Eurekahedge ILS Advisers index for specialist cat bond fund managers. Excess returns on independent variables are CRSP value weighted index (Equities), Bloomberg Barclays U.S. Corporate High Yield Total Return Index (High-yield bonds), Bloomberg Barclays US MBS Total Return Index (MBS), and carry trade (constructed by Lustig, Roussanov, and Verdelhan, 2011). Sample period is from 1/2006 to 12/2018, except for carry trade where sample ends in 10/2018. bootstrapped standard errors in parentheses. *, **, and *** denote statistical significance at 10%, 5%, and 1% level, respectively.

Coefficient	(1)	(2)	(3)	(4)	(5)
Intercept	0.272^{**}	0.258^{*}	0.266^{*}	0.319^{**}	0.268^{*}
	(0.136)	(0.139)	(0.158)	(0.131)	(0.160)
Equities	0.027				0.005
	(0.021)				(0.034)
High-yield bonds		0.057^{*}			0.038
		(0.031)			(0.051)
MBS			0.091		0.139
			(0.173)		(0.165)
Carry trade				0.031	0.008
				(0.040)	(0.049)
N	156	156	156	154	154
R^2	0.013	0.026	0.005	0.006	0.034

Table 1.7: Pricing of catastrophe market risk in subsamples

The table presents estimation results for subsamples of cat bonds using quarterly Fama-MacBeth regressions of the form

$$E_t(R^e_{i,t+1}) = \lambda_{0,t} + \lambda_{cat,t}\hat{\beta}_{i,t} + \varepsilon_{i,t} \,,$$

where the dependent variable is the expected excess return on cat bond i in the end of quarter t, and $\hat{\beta}_{i,t}$ is the simulated beta estimate. Row $\lambda_{cat,t} - E_t(R^e_{cat,t+1})$ reports the difference between the premium estimate and its theoretical value: excess return on cat bond market portfolio. R^2 is the time series average of cross-sectional figures. All prices are calculated using sheet prices. "Main" specification contains full sample results from Table 1.3. "Noncallables" and "Earthquake" columns include only non-callable bonds and bonds exposed to earthquake risk, respectively. "Parametric" column includes only bonds that have parametric of modeled loss as their trigger type. *, **, and *** denote statistical significance at 10%, 5%, and 1% level, respectively.

	Main	Noncallable	Earthquake	Parametric
$\lambda_{0,t}$	1.23***	1.20***	1.50***	1.28***
	(0.13)	(0.13)	(0.18)	(0.15)
$\lambda_{cat,t}$	2.06***	2.11***	1.98***	1.94***
	(0.18)	(0.18)	(0.30)	(0.25)
$\lambda_{cat,t} - E_t\left(R^e_{cat,t+1}\right)$	−1.10***	−1.05***	−1.18***	−1.23***
	(0.12)	(0.12)	(0.30)	(0.19)
N	63	63	63	63
R^2	0.49	0.51	0.43	0.49

Table 1.8: Price of non-extreme catastrope risk

The table presents estimation results for the regression

$$\Delta s_{i,j} = \lambda \Delta el_{i,j} + \varepsilon_{i,j} \,,$$

where the dependent variable is the difference in spreads between two consecutive tranches (i) of a catastrophe bond issue j. The regressor is the difference in log expected losses. Standard errors are in parentheses, and are clustered by issue j. Min(el_A) denotes the sample selection criteria for a given regression. For example, Min(el_A)=1.00 implies that only those pairs of cat bond trances where the safer trance has at least 1% annual expected loss are included in the sample. *, **, and *** denote statistical significance at 10%, 5%, and 1% level against a null hypothesis that $\lambda = 1$.

	(1)	(2)	(3)	(4)	(5)	(6)
λ	1.46***	1.42***	1.48***	1.44***	1.44***	1.49**
	(0.09)	(0.09)	(0.12)	(0.13)	(0.15)	(0.19)
N	151	125	92	57	45	34
N Clusters	106	89	67	42	33	25
Largest Cluster (%)	2.65	3.20	4.35	7.02	8.89	11.76
R^2	0.82	0.84	0.85	0.86	0.86	0.87
Min(el_A)	0.00	0.50	1.00	1.50	2.00	2.50

Table 1.9: Pricing of catastrophe market risk and bond-specific skewness

The table presents estimation results for quarterly Fama-MacBeth regressions of the form

$$E_t(R^e_{i,t+1}) = \lambda_{0,t} + \lambda_{cat,t}\hat{\beta}_{i,t} + \lambda_{skew,t} Skew_{i,t} + \varepsilon_{i,t}\,,$$

where the dependent variable is the expected excess return on cat bond i in the end of quarter t. $\hat{\beta}_{i,t}$ is the simulated beta estimate, and $Skew_{i,t}$ is the bond-specific estimate for return skewness, obtained from Eq. (1.11). R^2 is the time series average of cross-sectional figures. All prices are calculated using sheet prices. *, **, and *** denote statistical significance at 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)
λ_0	1.23***	5.33***	1.54***
	(0.13)	(0.43)	(0.28)
λ_{cat}	2.06***		2.03***
	(0.18)		(0.19)
λ_{skew}		0.20***	0.03*
		(0.03)	(0.01)
N	63	63	63
R^2	0.49	0.19	0.52

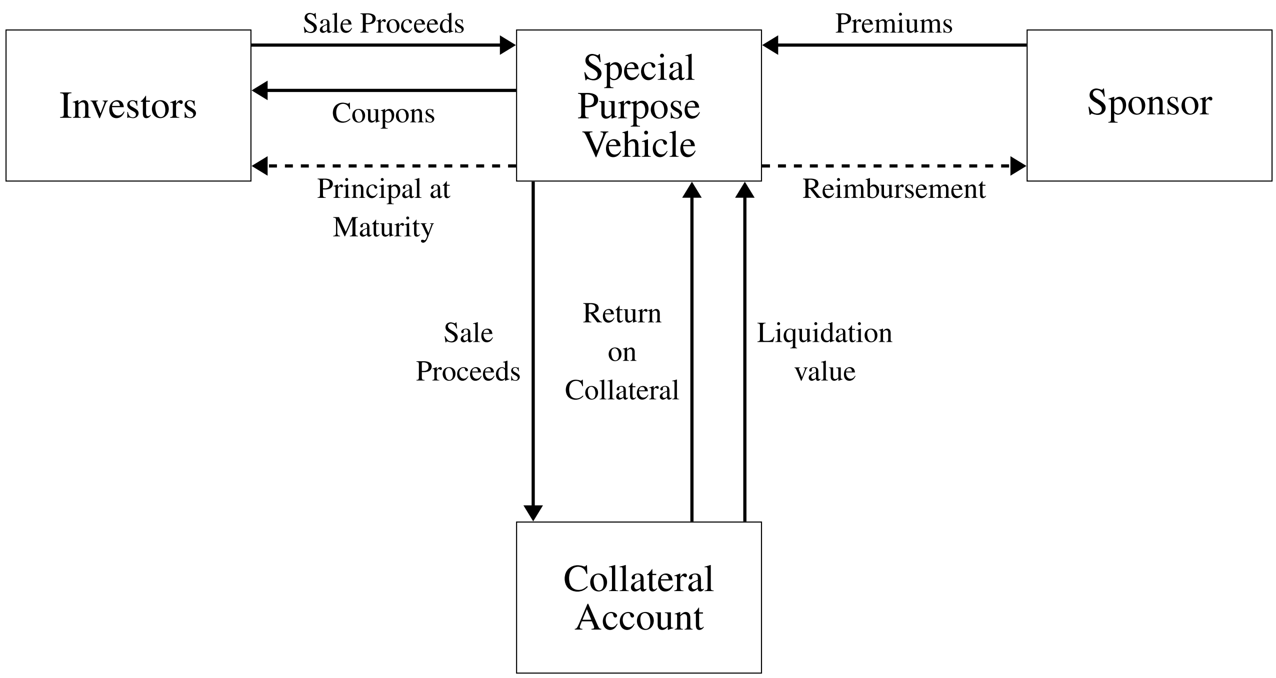

Figure 1.2: Typical cat bond deal structure

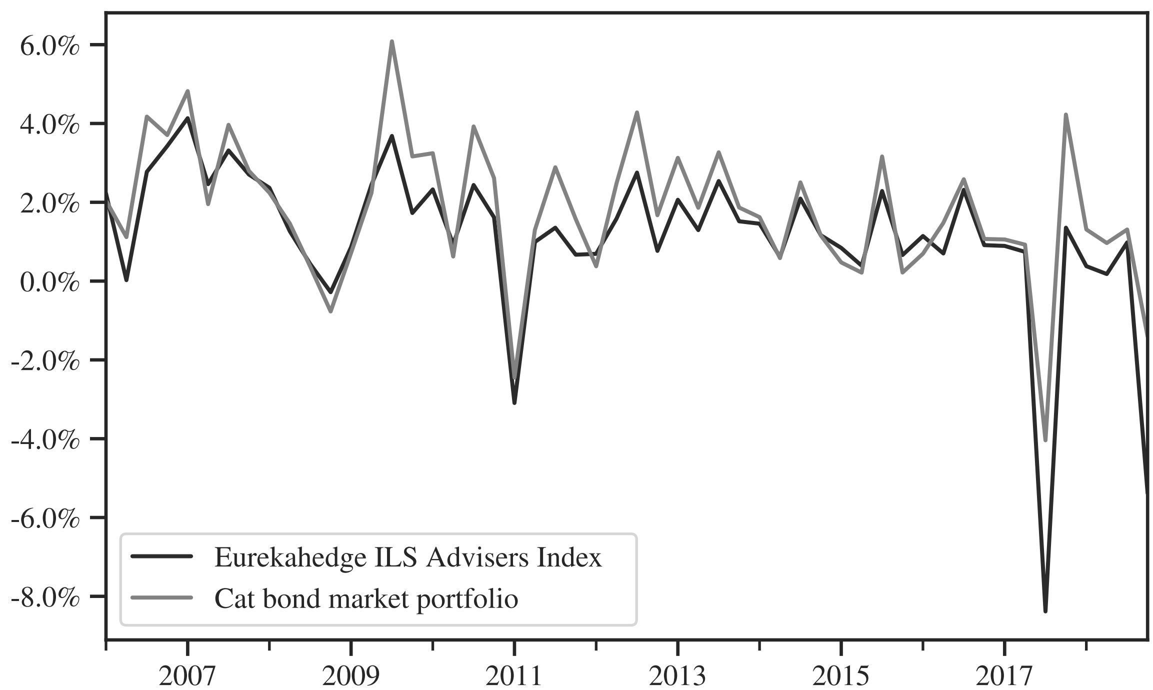

Figure 1.3: Returns on specialist funds and value-weighted portfolio of sample bonds

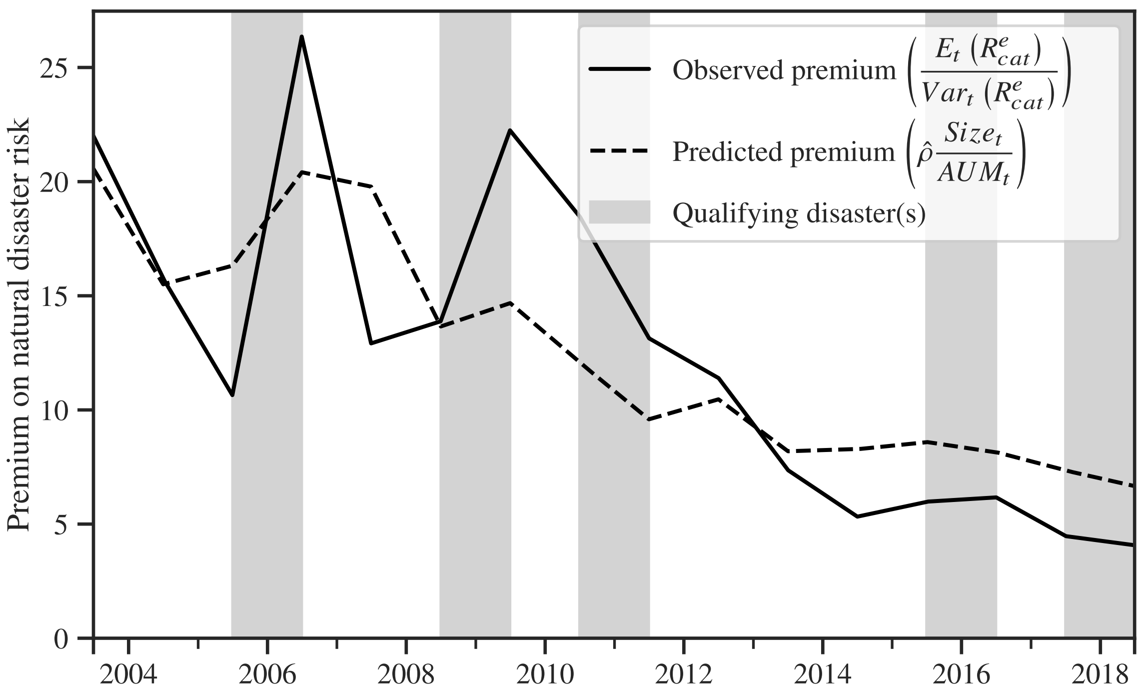

Figure 1.4: Time series evolution of price of natural disaster risk.

Description: The figure plots observed and predicted premium on natural disaster risk, where coefficient of relative risk aversion $\hat{\rho}$=6.1. R^e_{cat} is return on cat bond market portfolio, $Size_t$ is the size of the market, and AUM_t is total assets under management of the specialist funds in year t. Shaded regions indicate years (starting in July) during which qualifying loss events occurred.

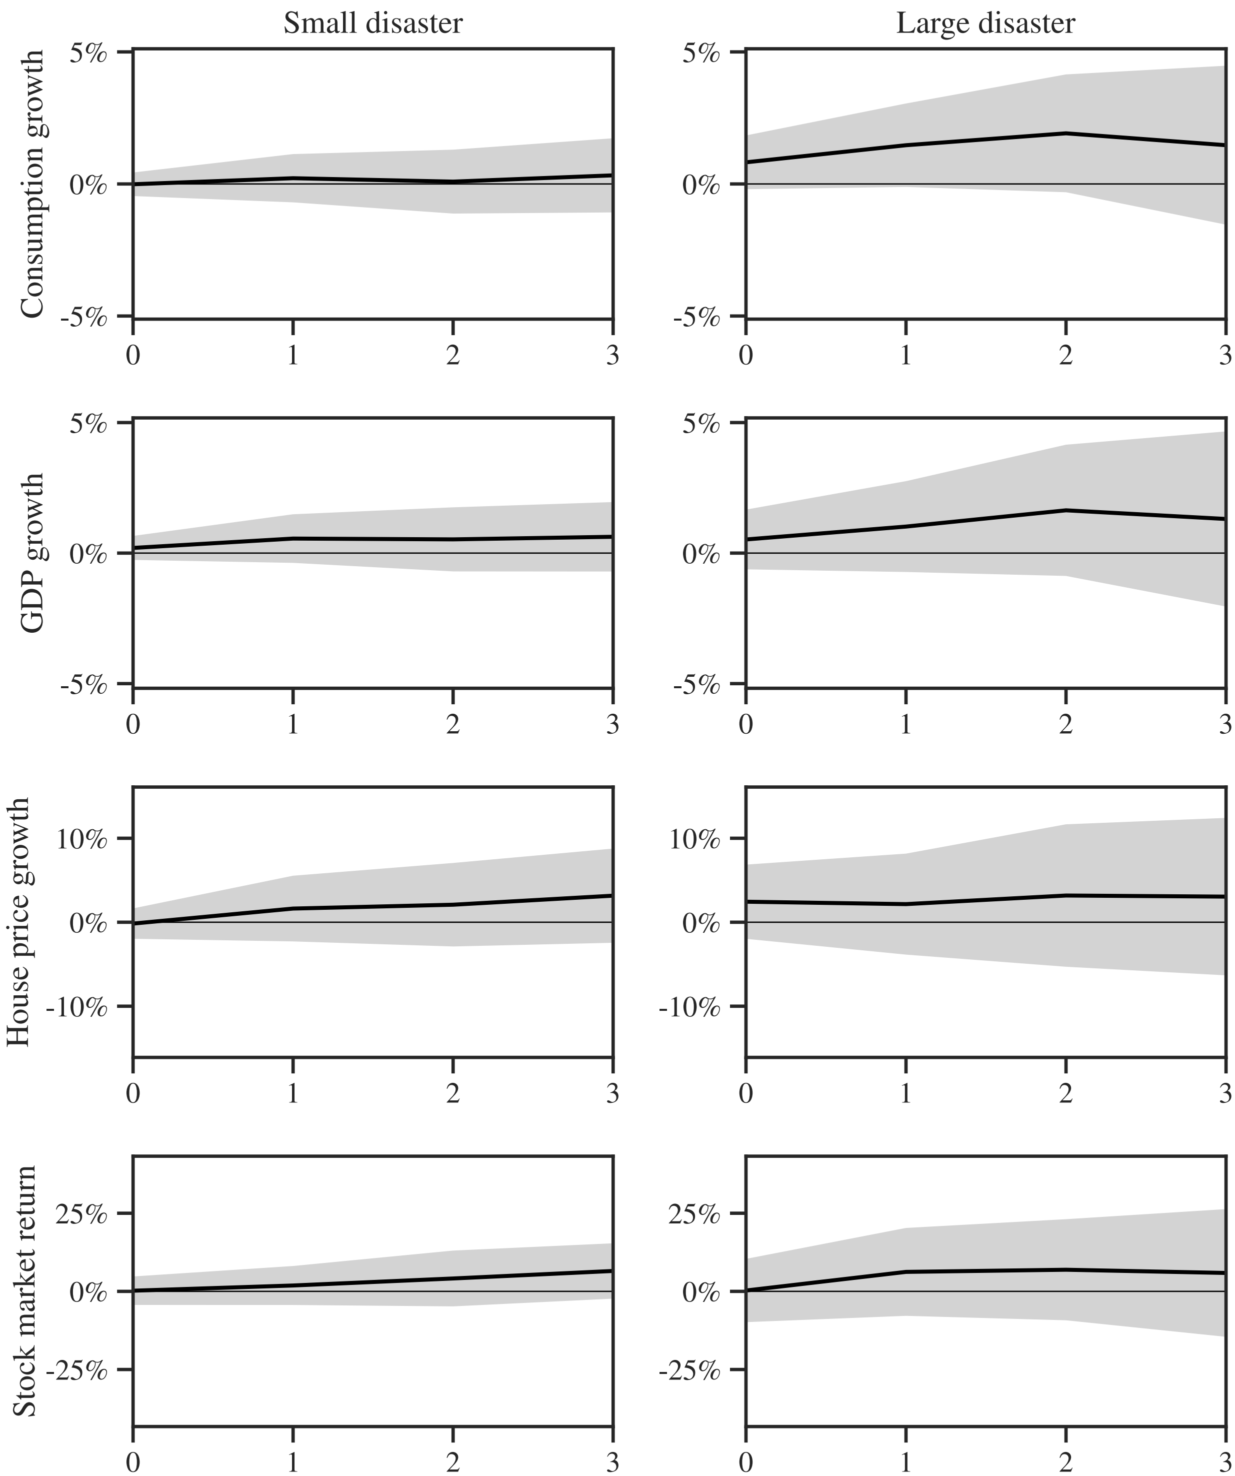

Figure 1.5: Impulse-responses of selected macroeconomic variables to natural disasters occurring in year 0.

Description: Shaded region shows 95% confidence interval on the estimates. Standard errors are clustered by year. Sample includes 13 developed countries from 1950 to 2016.

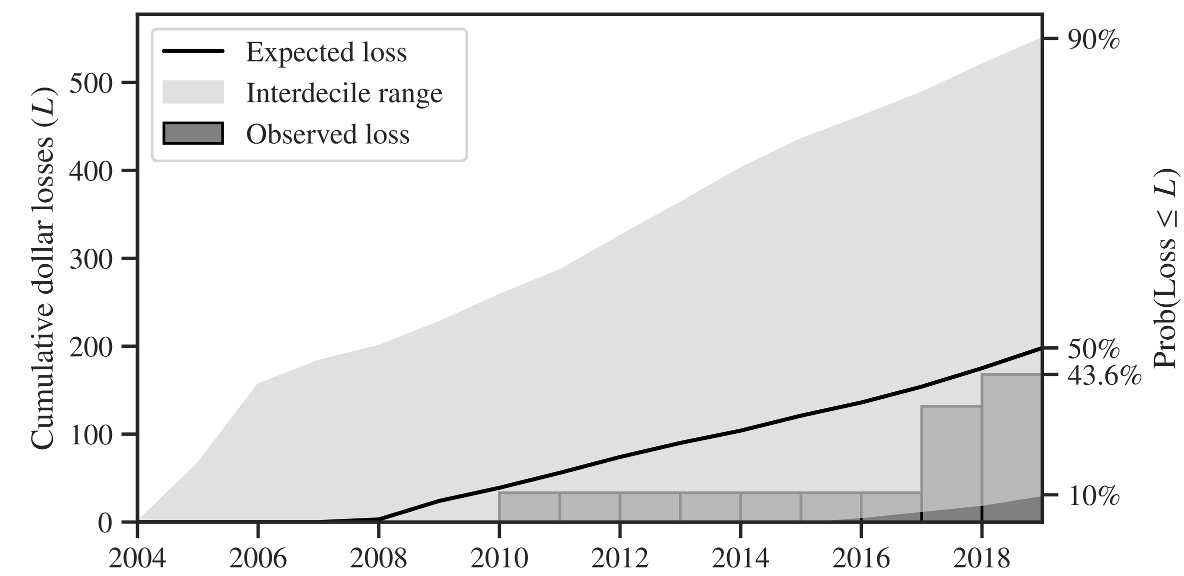

(a) Unbiased actuarial models

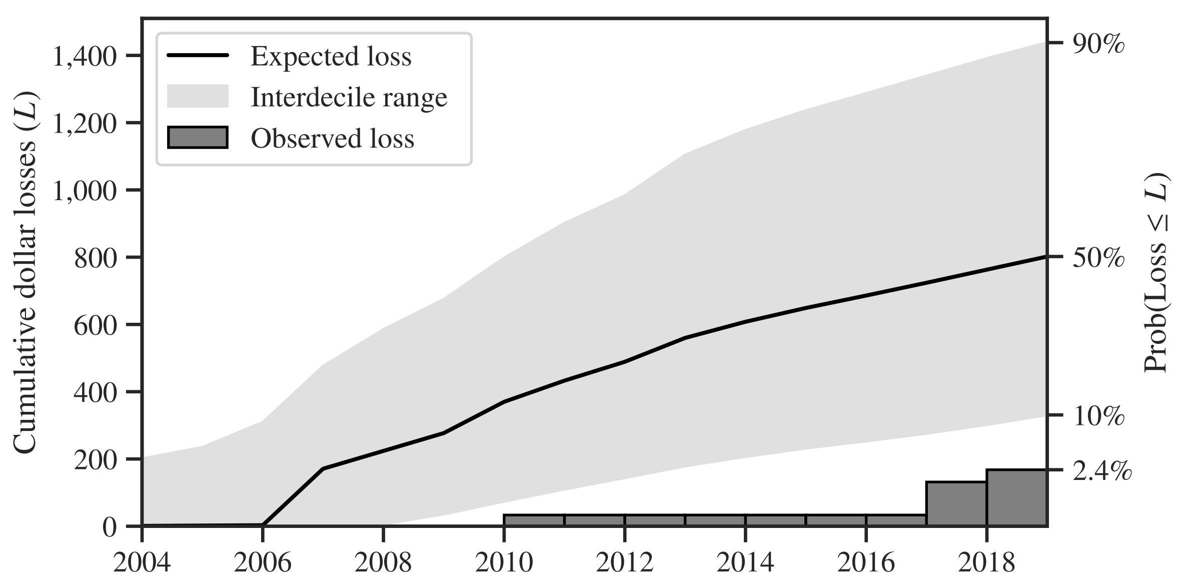

(b) Biased actuarial models

Figure 1.6: Predicted and actual losses on cat bonds due to natural disasters.

Description: Predicted losses calculated under the null that the actuarial models are unbiased (panel a), and under the null that the actuarial models are biased and the correct expected losses are implied by discount margins (panel b). The figures plot expected losses as a fraction of total capital invested from a strategy that invests in an equally weighted portfolio of single-peril cat bonds (rebalanced annually in the end of June). Actual cumulative losses are shown in bar graphs.

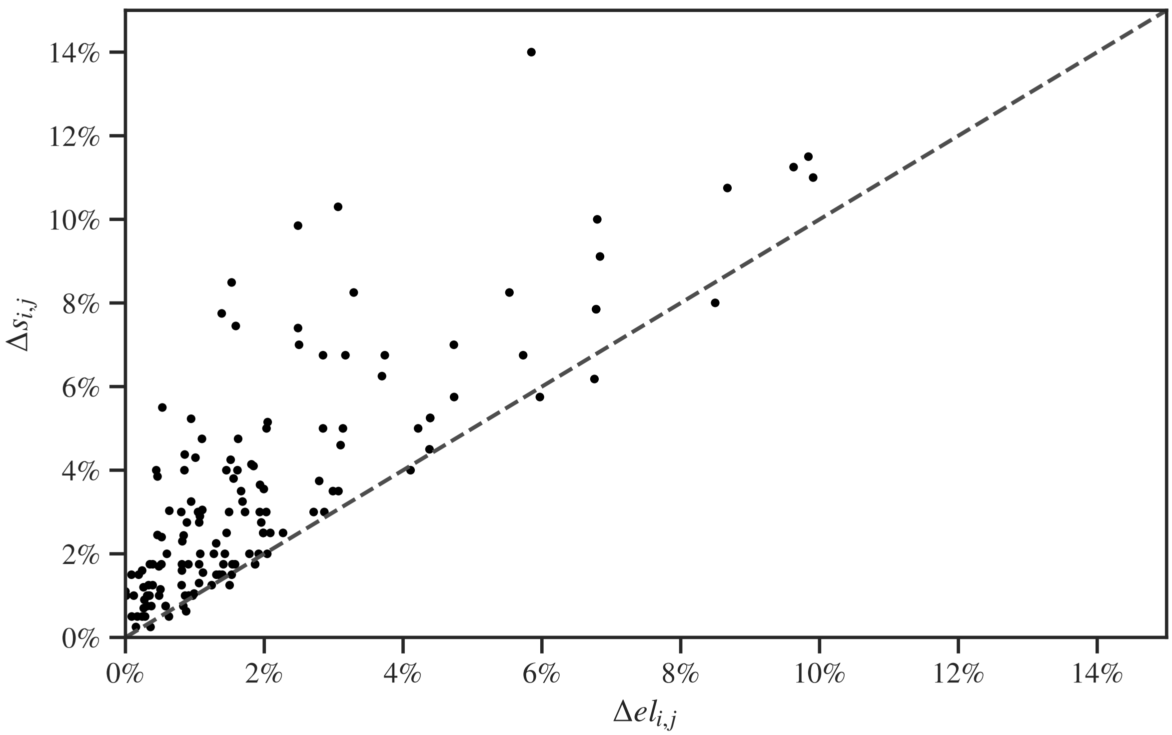

Figure 1.7: Differences in spreads and expected losses between two cat bond tranches

1.8 Appendix A: Data

1.8.1 Realized returns

Realized return on bond i in quarter $t+1$ is

$$R_{i,t+1} = \frac{P_{t+1} + C_{t+1}}{P_t}.$$

I start by constructing a quarterly data table that includes all bonds that are outstanding in a given quarter (adjusted for calls and maturity extensions). Then, I populate this table with prices using sheet prices (converted from discount margins). The resulting table for P_t suffers from a significant survivorship bias, because pricing information is often missing after a triggering event. I adjust for these cases using the following procedure. First, after a triggering event quarter, I use TRACE prices if available. Alternatively, I use sheet prices if TRACE prices are missing.[35] If a few rare cases, a price is missing for some particular quarter but is available in subsequent quarters. In such cases, I interpolate the missing value. Finally, if I don't observe any prices after a particular date, I set the remaining values equal to the estimated recovery value obtained from artemis.bm in September 2019.

Another reason for missing prices is that a bond was issued only recently (i.e., during the previous quarter). In such cases, I set the price equal to the issuance price. For some bonds, price information is also missing during the last quarter-end if the bond matures in the first month of the following quarter. In such cases, I set the price to par ($100) after verifying that the bond is not subject to any known losses. If the price is missing due to a call that has already been announced but not yet implemented, I set the price equal to the call price.

Finally, there are 11 sample bonds for which there are no sheet prices available or the availability is not consistent. Most of these bonds are associated with the early part of the sample (pre-2006)

[35] In a few cases, the observed prices are clearly inconsistent with the overall price sequence and ultimate recovery value. In these cases, I treat the prices as missing and use recovery values. The cases are Caelus Re V Ltd. (Series 2017-1) in Q3/2018, Carillon Ltd. Class A-1 in in Q1/2009 and Q4/2009, Gator Re Ltd. in Q1/2017 and Q3/2017, Mariah Re in Q4/2011, and Willow Re in Q4/2008 (TRACE).

and are hence not used in this study, or the prices are missing for all observations.

Next, I calculate P_{t+1} by adjusting for any delistings. First, I set the delisting price of any triggered bond equal to its estimated recovery value. Then, I set the delisting price of called bonds equal to the call price. Finally, I set the delisting price of all other bonds equal to the par value.

1.9 Appendix B: Proof of main predictions

Manager f's optimization problem can be written as:

$$\max_{\mathbf{q}'_f} \mathbf{q}'_f \left(E_0\left(\mathbf{P}_1\right) + \mathbf{C}\right) + \left(\frac{w_{m,0}}{\alpha} - \mathbf{q}'_f \mathbf{1}\right)(1+r) - \frac{\alpha\rho}{2}\mathbf{q}'_f \mathbf{\Omega} \mathbf{q}_f \, ,$$

where $\mathbf{q}_f$, $\mathbf{P}_1$ and $\mathbf{C}$ are $N \times 1$ vectors of quantities bought, prices, and coupons. $\mathbf{\Omega}$ is $N \times N$ variance-covariance matrix of $\mathbf{P}_1$.

First-order conditions are given by:

$$\underbrace{E_0\left(\mathbf{P}_1\right) + \mathbf{C} - \mathbf{1}(1+r)}_{E_0(\mathbf{R}^e)} = \alpha\rho\mathbf{\Omega}\mathbf{q}_f \, .$$

Market clearing ($Fq_i = \theta_i$, $\forall i$) implies that

$$E_0\left(R_i^e\right) = \frac{\alpha\rho}{F} Cov_0\left(P_{i,1}, P_{cat,1}\right) \, ,$$

where $P_{cat,1} = \mathbf{P_1}\theta$. Note that by definition, $P_{i,0} = 1 \, \forall i$. As a result,

$$E_0\left(R_i^e\right) = \alpha a \rho Cov_0\left(R_i^e, R_{cat}^e\right) \, ,$$

where $a = \sum_i \theta_i / F$ is the average position that a fund has in the cat bond markets. Since the previous expression holds also for the market portfolio,

$$E_0\left(R_i^e\right) = \underbrace{\frac{Cov_0\left(R_i^e, R_{cat}^e\right)}{Var_0\left(R_{cat}^e\right)}}_{\beta_i} E_0\left(R_{cat}^e\right) ,$$

and

$$E_0\left(R_{cat}^e\right) = \alpha a \rho Var_0\left(R_{cat}^e\right) .$$

Note that since the $E_0\left(R_{cat}^e\right) > 0$, constraint in Eq. (1.6) is always binding. This is because $E_0\left(M_1 R_{cat}^e\right) > 0$ and as a result the outside investors are willing to give the managers the maximum amount of capital possible.

1.10 Appendix C: A model with priced peso states

Let $x_i \sim U(0,1)$. x_i is a random variable that represents the severity of a particular set of disasters, with smaller values of x_i indicating a more severe disasters. If the law of one price holds, we have an SDF of the form $M\left(x_i, \mathbf{z}\right)$, where $\mathbf{z}$ represents a vector of relevant macroeconomic variables that affect pricing. Without loss of generality, x_i and $\mathbf{z}$ are independent.

Consider two cat bonds, A and B, whose payoffs X_A and X_B are given by Eq. (1.2), and that are exposed to the same perils. Let the exhaustion probability and expected loss of bond A to be smaller than those of bond B ($\underline{x}_A \leq \underline{x}_B$ and $el_A < el_B$) implying that bond A is safer than bond B.[36]

Now, let us make two assumptions.

Assumption A1 (Small disasters are not priced):

$$M\left(x_i, \mathbf{z}\right) = M\left(\mathbf{z}\right),\ x_i \geq x_i^* . \tag{1.19}$$

Assumption A1 implies that for a natural disaster c, there is some "rare event threshold" x_c^* such

[36] In actual cat bond issuances, the payoff of A typically dominates the payoff of B for all $\mathbf{x}$, but we only need these weaker conditions.

that risks related to any disaster that is smaller than this threshold are diversifiable and hence not priced. Here, "small disasters" should be interpreted broadly: they can cause large local damages, but are not big enough to cause economy-wide disruptions. Generally speaking, the natural disasters we have observed in the panel of 13 developed economies since 1950 are likely to fall into this category, because these events have not been associated with large drops in GDP or aggregate consumption.

Assumption A2 (Safe bond exhausts before "rare event threshold" is hit):

$$\underline{x}_A \geq x_A^* . \tag{1.20}$$

Assumption A2 will later impose a restriction on the sample, but is otherwise relatively weak. It limits the predictions that follow to concern only those cat bonds that are risky enough that even the small disasters can cause them to trigger. Put differently, Assumption A2 requires that every bond in our sample is risky enough that they would lose all principal in any priced peso state.

Figure 1.15 illustrates Assumptions A1 and A2 graphically. Assumption A1 requires that states where the scale of a disaster is below some threshold x^* are not priced. The benchmark model discussed in Section 1.2.2 is a special case of this setting in which x_A^* is arbitrarily small (i.e., every disaster is below the pricing threshold). Assumption A2 requires that by the time we hit x_A^*, both the risky and the safe bonds have fully lost their principal, which follows from $\underline{x}_B \geq \underline{x}_A \geq x_A^*$. As a consequence of these assumptions, the payoffs of Bonds A and B differ only in states of the world that are not associated with any aggregate risk.

The price of bond i is given by:

$$
\begin{aligned}
P_{i,0} &= E_0\left(M\left(x_i, \mathbf{z}\right) X_{i,1}\right) \\
&= \left(1 - \underline{\mathrm{p}}_i\right) E_0\left(M\left(x_i, \mathbf{z}\right) X_{i,1} | x_i > \underline{\mathrm{x}}_i\right) \\
&= E_0\left(M\left(\mathbf{z}\right) | x_i > \underline{\mathrm{x}}_i\right) \left(1 - \underline{\mathrm{p}}_i\right) E_0\left(X_i | x_i > \underline{\mathrm{x}}_i\right) \\
&= E_0\left(M\left(\mathbf{z}\right) | x_i > \underline{\mathrm{x}}_i\right) E_0\left(X_i\right)
\end{aligned}
$$

where the third equality is due to Assumption A1. Hence, the continuously compounded yield y_i is given by:

$$
y_{i,0} = -\ln E_0\left(M\left(\mathbf{z}\right) | x_i > \underline{\mathrm{x}}_i\right) - \ln E_0\left(X_i\right),
$$

Now, since $y_i = r + s_i$ and $E_0\left(X_i\right) = (1 - el_i)$, we have that:

$$
\Delta s = \Delta el\,,
$$

where $\Delta s = s_B - s_A$ is the difference in continuously compounded spreads between the risky and the safe cat bond, and $\Delta el = \ln(1 - el_A) - \ln(1 - el_B)$.

1.11 Appendix D: Tables and figures

Table 1.10: An example to illustrate variation in betas and expected returns

The table presents information on five selected bonds in June 2013. Risk categories are North Atlantic Hurricane (NAH) and Pacific Hurricane (PH). True exposures (which are not observed systematically for all bonds) are U.S. East and Gulf coasts for the first three bonds, Pacific coast of Mexico for the fourth bond, and Atlantic coast of Mexico for the fifth bond. el_i indicate expected actuarial losses. $\hat{\beta}_{i,t}$ are simulated beta estimates that are functions of Category and el_i. $\hat{\beta}_{i,t} E_t\left(R^e_{cat}\right)$ is predicted premia, and $E_t\left(R^e_i\right)$ observed premia. Premium for cat bond market portfolio $E_t\left(R^e_{cat}\right)$ is 3.19%. Full names of included bonds are Mythen Ltd. 2012-1 Class E, Mythen Ltd. 2012-1 Class A, Mythen Re Ltd. 2012-2 Class C, MultiCat Mexico Ltd. 2012-1 Class C, and MultiCat Mexico Ltd. 2012-1 Class B.

Bond	Category	True exposure	el_i	$\hat{\beta}_{i,t}$	$\hat{\beta}_{i,t} E_t\left(R^e_{cat}\right)$	$E_t\left(R^e_i\right)$
Mythen E	NAH	U.S. Atlantic	0.8%	0.96	3.1%	3.2%
Mythen A	NAH	U.S. Atlantic	1.1%	1.24	4.0%	3.7%
Mythen C	NAH	U.S. Atlantic	3.8%	2.40	7.6%	8.1%
MultiCat Mexico C	PH	Mexico Pacific	4.3%	0.14	0.4%	1.1%
MultiCat Mexico B	NAH	Mexico Atlantic	2.6%	2.21	7.0%	2.8%

Table 1.11: Pricing of catastrophe market risk

The table presents estimation results for cross-sectional regressions of the form

$$E_t(R^e_{i,t+1}) = \lambda_{0,t} + \lambda_{cat,t}\hat{\beta}_{i,t} + \varepsilon_{i,t} \,,$$

where the dependent variable is the expected excess return on cat bond i in the end of June of year t, and $\hat{\beta}_{i,t}$ is the simulated beta estimate. Standard errors are clustered by bond issue. Column $\lambda_{cat,t} - E_t(R^e_{cat,t+1})$ reports the difference between the premium estimate and its theoretical value: excess return on cat bond market portfolio. The last row contains quarterly time series averages of different parameter estimates, with associated Fama and MacBeth (1973) standard errors. R^2 in the last row is the time series average of cross-sectional figures. All prices are calculated using actual trade prices from TRACE. Bolded coefficients are significant at 1% level.

t	$\lambda_{0,t}$	(t-stat)	$\lambda_{cat,t}$	(t-stat)	$\lambda_{cat,t} - E_t\left(R^e_{cat,t+1}\right)$	(t-stat)	R^2	N	$N_{clusters}$
2005	0.76	1.75	**1.20**	3.39	-0.80	-2.26	0.40	16	12
2006	-1.93	-2.56	**7.57**	11.71	**1.99**	3.08	0.87	29	15
2007	**1.52**	5.32	**3.59**	9.58	**-1.32**	-3.53	0.87	23	17
2008	1.82	1.94	**2.41**	3.38	-1.35	-1.89	0.53	18	15
2009	**4.36**	4.21	**3.61**	3.26	**-3.60**	-3.24	0.53	19	15
2010	**3.47**	7.52	**1.40**	3.79	**-3.23**	-8.72	0.38	26	19
2011	1.29	1.48	**2.52**	3.51	-1.44	-2.00	0.60	18	12
2012	**1.97**	5.17	**4.14**	11.72	**-2.49**	-7.05	0.86	22	20
2013	**0.95**	4.98	**2.14**	8.40	**-1.26**	-4.95	0.76	30	26
2014	**1.36**	9.88	**0.96**	5.08	**-1.63**	-8.64	0.47	34	28
2015	**1.37**	7.47	**0.99**	6.37	**-1.41**	-9.03	0.55	35	27
2016	**0.98**	6.34	**0.91**	4.46	**-0.83**	-4.04	0.55	29	23
2017	**0.82**	3.27	0.77	2.42	-0.41	-1.29	0.14	35	27
2018	**0.72**	3.30	0.46	2.44	**-0.57**	-3.02	0.17	31	22
FM	**1.51**	9.32	**1.95**	10.02	**-1.40**	-9.47	0.46	57	

Table 1.12: Effect of natural disasters on selected macroeconomic variables

The table presents estimation results for Jordà (2005) local projections

$$\Delta_h y_{i,t+h} = \gamma_i + \gamma_t + b_1 Small_{i,t} + b_2 Large_{i,t} + \varepsilon_{i,t}\,,$$

where $\Delta_h y_{i,t+h} = y_{i,t+h} - y_{i,t-1}$ is $h+1$ year change in the macroeconomic variable of interest, that include real consumption growth per capita, real GDP growth per capita, nominal house price index growth and return on stock market index. $Small_{i,t}$ is an indicator variable that is equal to 1 if natural disasters caused economic damages between 0.1% and 1% of previous year's GDP in country i. $Large_{i,t}$ is an indicator variable that is equal to 1 if damages were larger than 1%. γ_i and γ_t are country and year fixed effects (FE), respectively. Sample period is 1950-2016 and includes 13 developed countries. Standard errors in parentheses are clustered by year. *, **, and *** denote statistical significance at 10%, 5%, and 1% level, respectively.

	Consumption growth				GDP growth			
	$h=0$		$h=2$		$h=0$		$h=2$	
$Small_{i,t}$	−0.011	−0.259	0.088	−1.067	0.198	0.056	0.526	−0.758
	(0.227)	(0.258)	(0.616)	(0.718)	(0.234)	(0.298)	(0.625)	(0.775)
$Large_{i,t}$	0.820	0.896	1.917*	2.598*	0.522	0.950	1.639	2.754
	(0.518)	(0.601)	(1.136)	(1.565)	(0.583)	(0.738)	(1.281)	(1.820)
FE	Yes	No	Yes	No	Yes	No	Yes	No
N	832	832	832	832	832	832	832	832
R^2	0.003	0.003	0.004	0.007	0.002	0.003	0.004	0.006

	House price growth				Stock market return			
	$h=0$		$h=2$		$h=0$		$h=2$	
$Small_{i,t}$	−0.161	−0.395	2.086	−0.529	0.215	0.571	4.126	5.020
	(0.921)	(0.916)	(2.532)	(2.667)	(2.320)	(2.873)	(4.547)	(5.551)
$Large_{i,t}$	2.444	3.093	3.180	1.834	0.275	−3.811	6.921	7.875
	(2.250)	(1.971)	(4.319)	(4.486)	(5.155)	(5.744)	(8.239)	(8.110)
FE	Yes	No	Yes	No	Yes	No	Yes	No
N	743	743	743	743	832	832	832	832
R^2	0.002	0.003	0.002	0.000	0.000	0.001	0.002	0.002

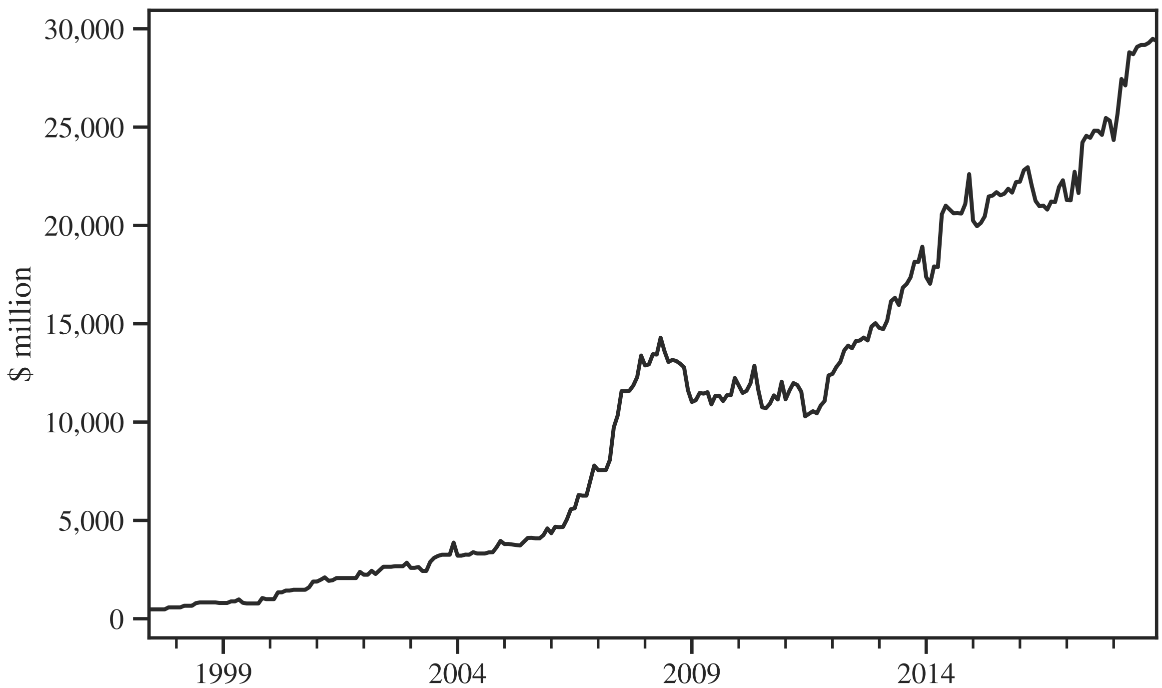

Figure 1.8: Total face value outstanding for sample bonds

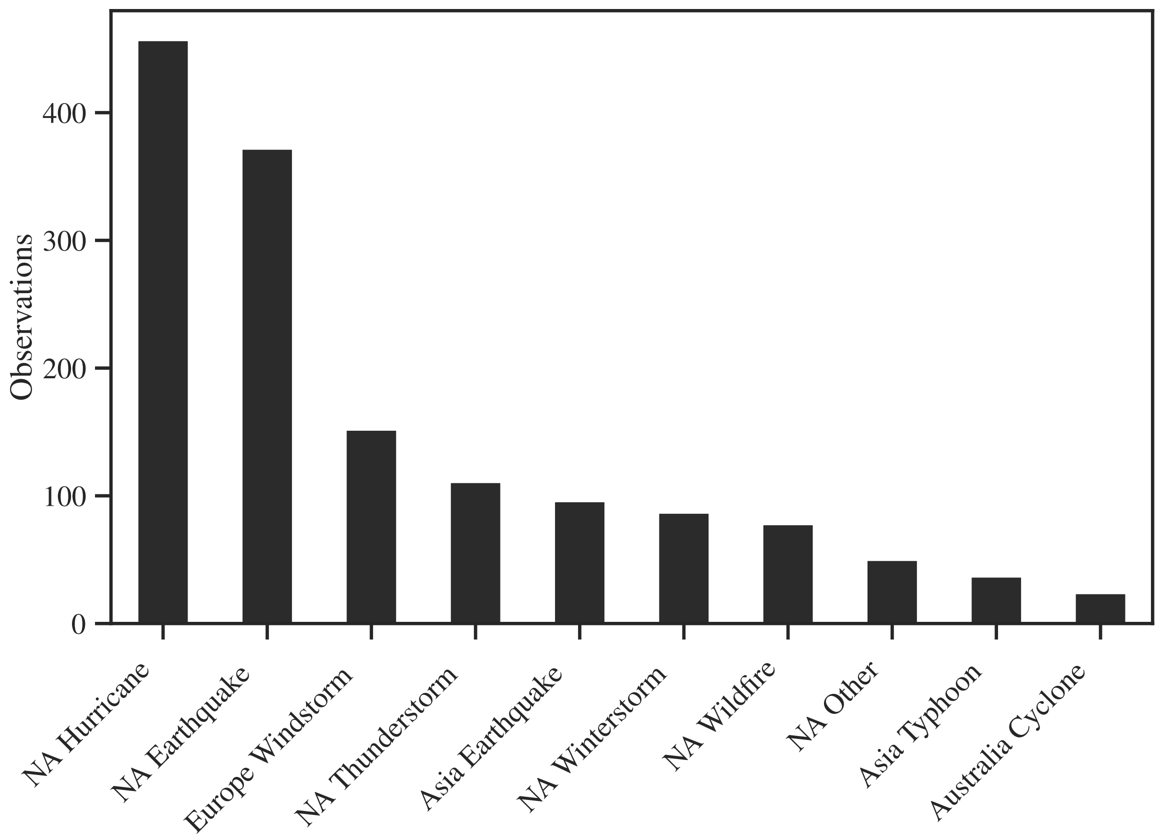

Figure 1.9: Ten most covered perils

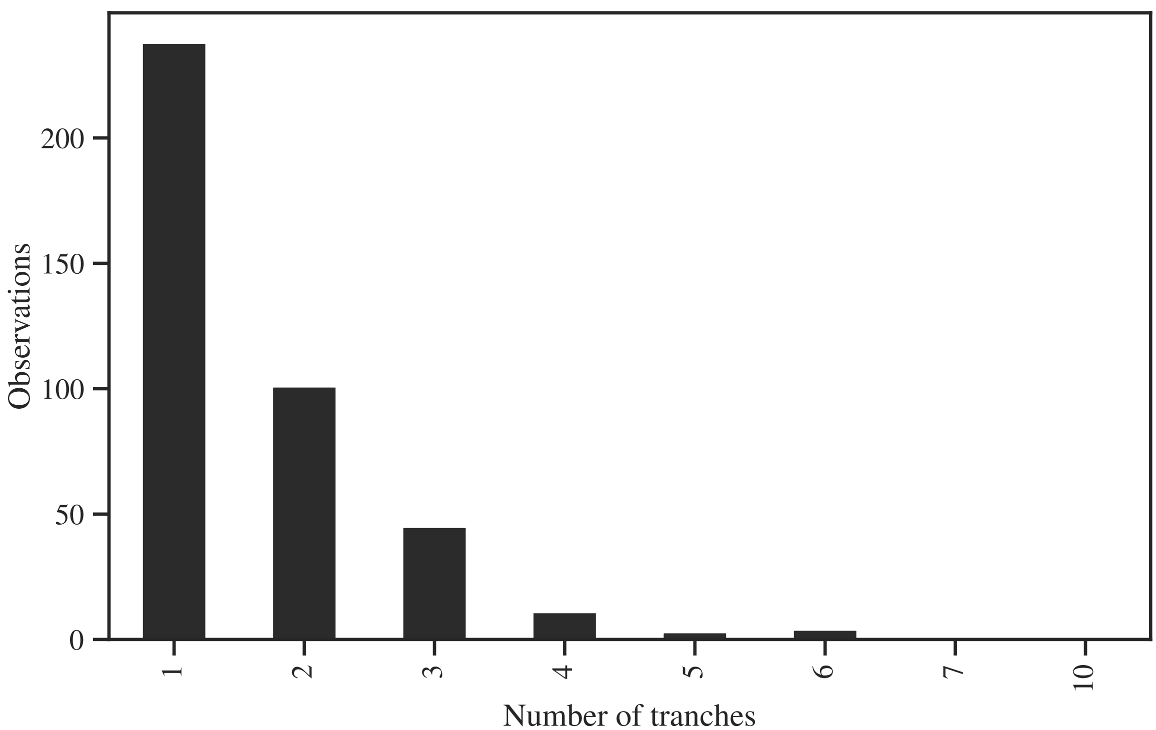

Figure 1.10: Distribution of number of tranches per issuance

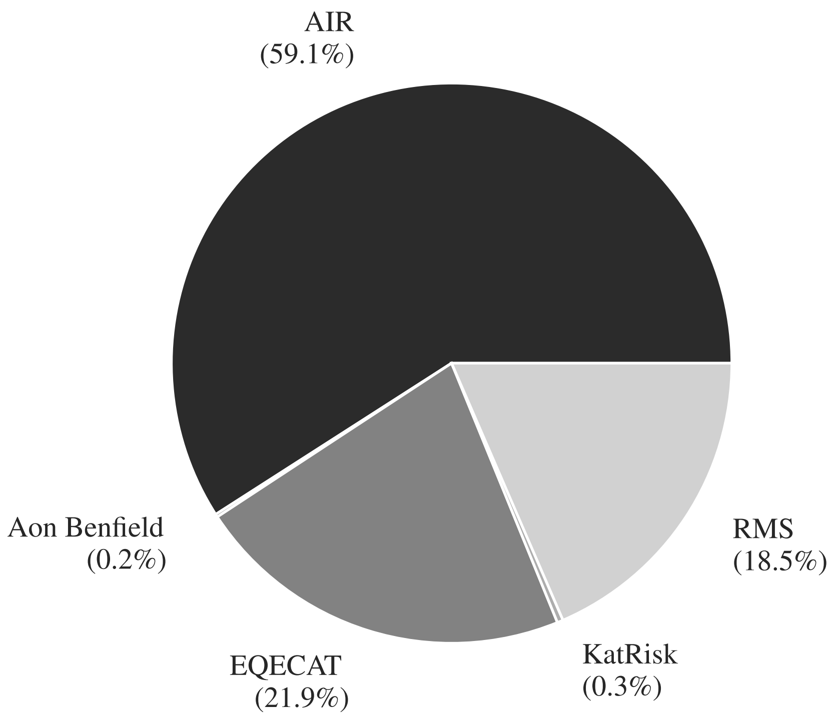

Figure 1.11: Market share of risk modeling companies

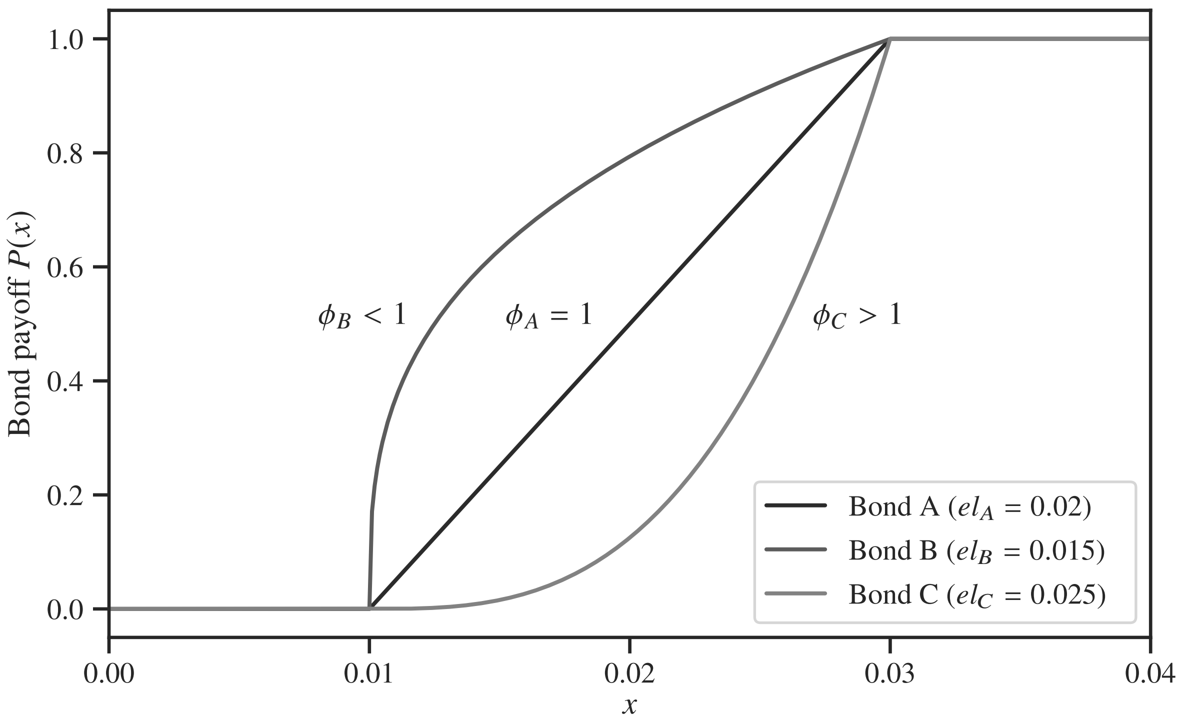

Figure 1.12: Illustration of loss function P

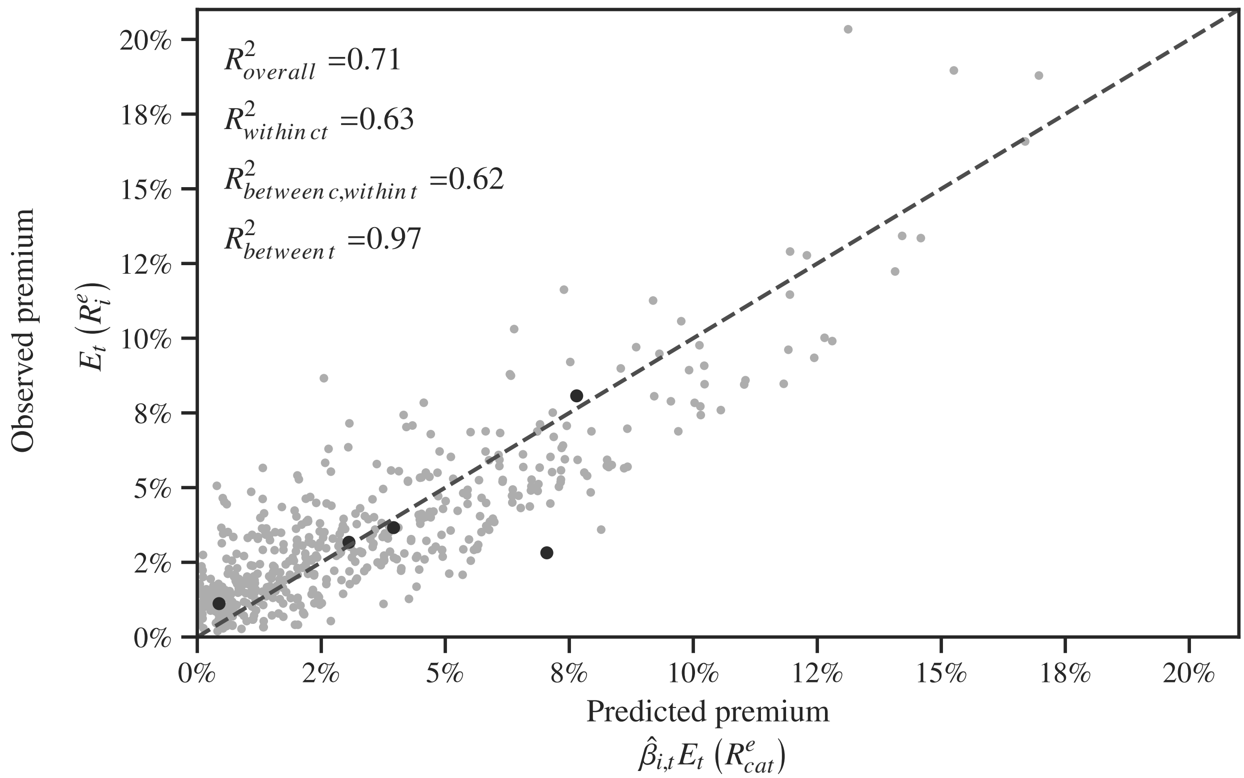

Figure 1.13: Observed vs. predicted risk premium.

Description: The figure plots expected excess returns (discount margin - expected actuarial loss) on sample cat bonds between 2003 and 2018 against their predicted values from a single-factor intermediary model. Bonds' betas are estimated from 500,000 years of simulated disaster data, and each dot represents a single cat bond observation in the end of June of a given year. Five highlighted observations are the ones discussed in Table 1.10. R-squared measures are as follows: $R^2_{overall}$ is pooled, $R^2_{within\,ct}$ is within year and risk category, $R^2_{between\,c,within\,t}$ is between risk categories and within a year, and $R^2_{between\,t}$ is between years.

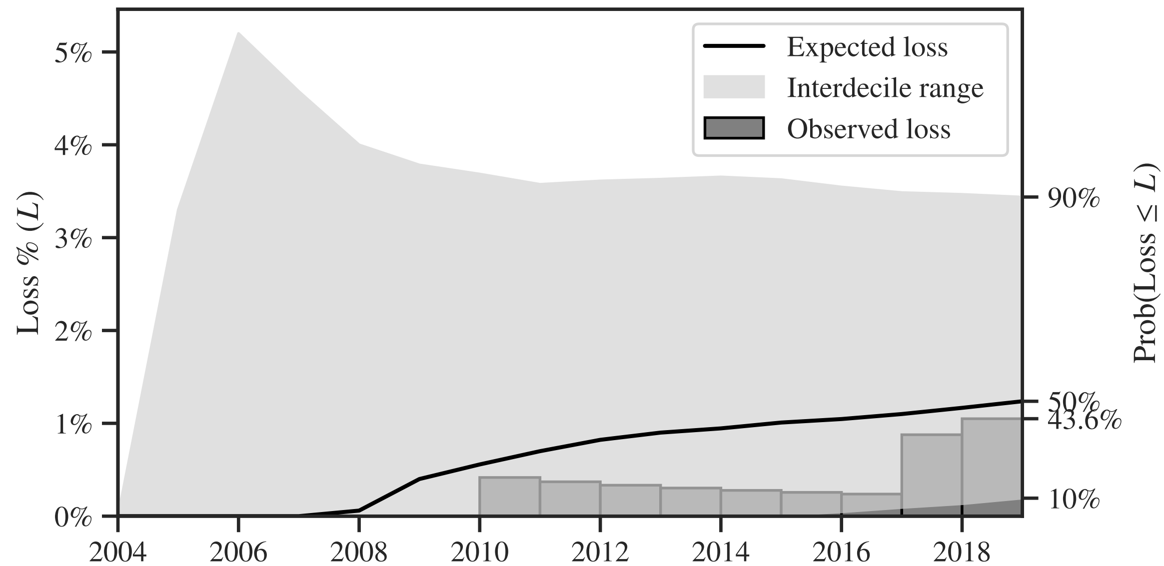

(a) Unbiased actuarial models

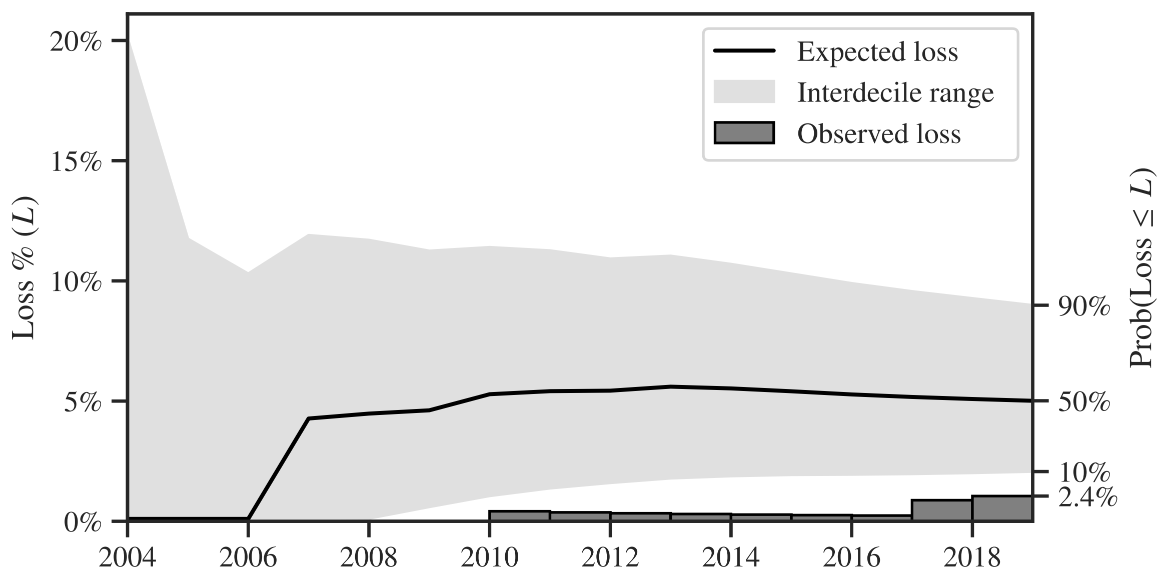

(b) Biased actuarial models

Figure 1.14: Predicted and actual losses on cat bonds due to natural disasters.

Description: Predicted losses calculated under the null that the actuarial models are unbiased (panel a), and under the null that the actuarial models are biased and the correct expected losses are implied by discount margins (panel b). The figures plot losses as a fraction of total capital invested from a strategy that invests in an equally weighted portfolio of single-peril cat bonds (rebalanced annually in the end of June). Actual cumulative losses are shown in bar graphs.

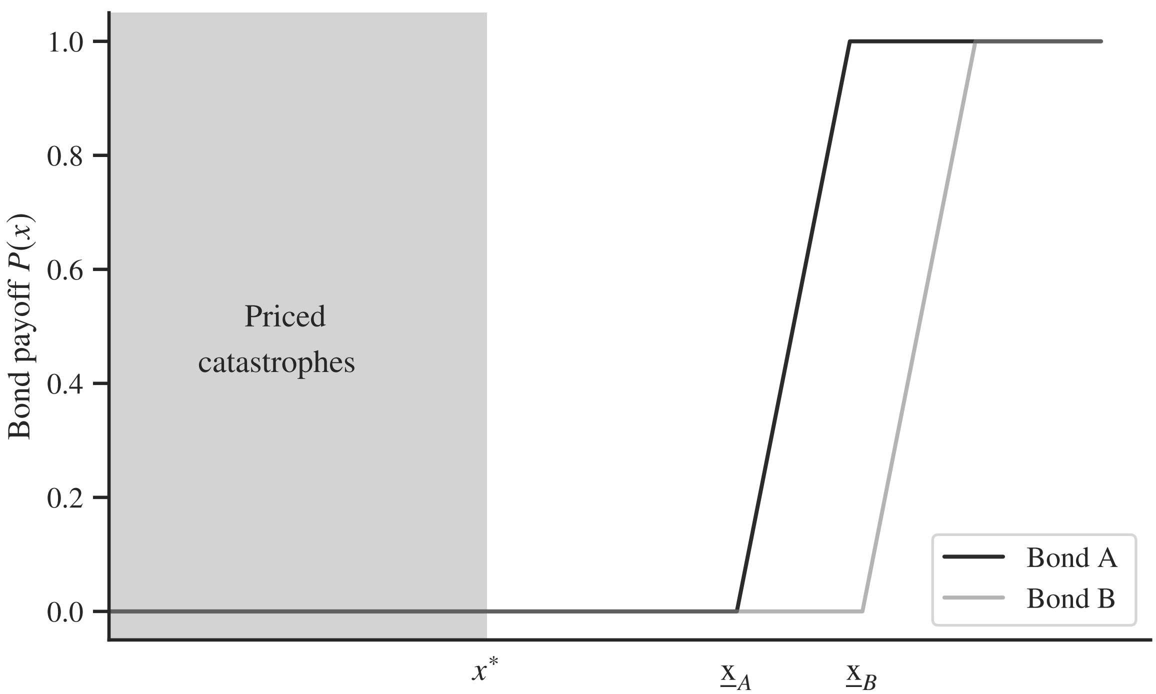

Figure 1.15: Illustration of Assumptions A1 and A2

Chapter 2: Taking the Cochrane-Piazzesi Term Structure Model Out of Sample: More Data, Additional Currencies, and FX Implications

2.1 Introduction

In two seminal papers, Cochrane and Piazzesi (2005, 2008) document a strong one-factor structure in the unconstrained predictability of one-year-ahead excess returns on U.S. dollar zero-coupon bonds of several maturities. The first principal component of the four unconstrained expected excess returns explains 99.7% of the variance of these expected returns. This finding motivates a constrained model for which Cochrane and Piazzesi (2005) note (p. 142), "The same function of forward rates forecasts holding period returns at all maturities. Longer maturities just have greater loadings on this same function." They model this constrained system with a two-step estimation in which they first estimate the forecasting factor, which is labeled the 'CP factor' in much of the subsequent literature, by regressing the average future annual excess rates of return on two, three, four, and five year bonds onto a set of forward rates or forward spreads. Then, they regress each excess return on the CP forecasting factor to get the factor loadings. The constrained model fits the unconstrained expected excess return data remarkably well. They also demonstrate that their bond market forecasting factor predicts excess returns in the U.S. stock market, which strengthens the case that it is capturing risk premiums. Cochrane and Piazzesi (2008) reverse engineer an affine term structure model (ATSM) that has the forecasting properties uncovered in the constrained regressions.

This paper examines whether analogous one-factor forecasting structures exist in the predictability of the excess returns on zero-coupon bonds denominated in other currencies, and we find that they do. For all currencies in the G10 except the New Zealand dollar, we find that the first principal component of the four unconstrained expected returns explains over 99% of the variance in expected returns in our first sample. We initially examine samples that end in 2003, the end of the sample in Cochrane and Piazzesi (2005). For the constrained model, we find that the factor loadings are quite similar across currencies, ranging between 0.37 and 0.47 for the two-year bond, between 0.80 and 0.87 for the three-year bond, and between 1.19 and 1.23 for the four-year bond. In contrast to this similarity, we find the parameter estimates on the forward spreads that generate the CP factors to be quite dissimilar across the different currencies. These estimates could reflect differences in the ultimate sources of risk premiums due to alternative monetary and fiscal policies and other macroeconomic shocks across the countries, they could reflect short samples with a corresponding failure of rational expectations econometrics, or they could be the result of multicollinearity, which seems particularly likely in some countries where we see large coefficients of approximately equal and opposite signs.

We then examine data from 2004-2016, and we again find strong one-factor forecasting structures in all nine currencies with the first principal component of the four unconstrained expected returns explaining over 95% of the variance in the unconstrained expected returns in each case. While the factor loadings in the second sample are quite similar to those of the first sample, ranging between 0.28 and 0.38 for the two-year bond, between 0.73 and 0.82 for the three-year bond, and between 1.22 and 1.29 for the four-year bond, the estimates of the slope coefficients on the CP factors change quite a bit. Consequently, we generally reject the hypothesis of equality of the coefficients in the CP factors across the two samples. For example, the coefficient estimate on the U.S. two-year forward spread changes from 2.32, with a standard error of (1.63), to -7.57 (2.62), while the coefficient on the Japanese two-year forward spread changes from -14.99 (3.75) to 3.33 (1.73).

Because foreign exchange rates and the term structures of interest rates in the two currencies are closely linked in theory to the stochastic discount factors of the two currencies, we derive new predictions from the Cochrane and Piazzesi (2008) ATSM for the predictability of excess rates of return on uncovered foreign currency investments. We find that the CP factors from the bond markets of the two currencies and their squared values should forecast the excess rate of return on uncovered foreign currency investments between the two currencies. Because tests of uncovered interest rate parity generally find that interest rate differentials predict excess currency returns, which is not true in our theoretical derivation of currency risk premiums from the affine model, we investigate whether the CP factors of the two currencies drive out the interest rate differentials in predicting excess currency returns. We find that they do not as the interest differentials remain the only significant predictors. For example, in forecasting the excess return to investing the U.S. dollar in the British pound, we find a coefficient on the pound-dollar interest differential of 1.86 (0.89) while the coefficients on the USD and GBP CP factors are 0.24 (0.62) and -0.24 (1.20), respectively. In this first sample, the coefficients on the interest differentials are also mostly greater than one, which is consistent with the typical findings in the literature that rejects the uncovered interest rate parity hypothesis and motivates the carry trade. In this analysis, though, we also show substantial differences in estimated coefficients across our two samples with the coefficients on the interest differentials generally reversing sign in the second sample, which is consistent with the deterioration of the carry trade during and after the global financial crisis.

Although the Cochrane and Piazzesi (2005) model characterizes the bond market data quite well in any given sample, as noted above, we generally reject that the parameters are the sample across our two samples. As an additional investigation of the stability of the model, we explore the recursive out-of-sample predictions of the Cochrane and Piazzesi (2005) model. Recursive forecasts of excess rates of return in the bond markets from the estimated Cochrane and Piazzesi (2005) model are generally unable to beat the recursive forecasts from the historical averages of excess rates of return on bonds. While these findings are perhaps unsurprising given that the out-of-sample period contains the global financial crisis, they demonstrate the necessity of modeling risk premiums while allowing for structural change. We leave this challenging task for future research. For completeness, we also show that recursive estimation and forecasting from the empirical version of our international extension of the Cochrane and Piazzesi (2008) model similarly fails to beat historical average excess currency returns in out-of-sample tests. This is completely unsurprising given the in-sample failure of the CP factors to predict currency returns and the changes in the signs of the coefficients across the two samples.

The last sections of the paper relate our findings to the existing literature and discuss some alternative modeling approaches that may improve our understanding of the term structure of interest rates and the predictability of bond and currency returns.

2.2 The Cochrane-Piazzesi Term Structure Model

In presenting the model, we mostly adopt the notation of Cochrane and Piazzesi (2005). The presentation can be thought of as referring to the term structure of a generic currency. For simplicity, we suppress currency subscripts in laying out the basic term structure model.

The natural logarithm of the price of a pure discount bond at time t that matures in n years and pays one unit of currency at that time is denoted $p_t^{(n)}$. The time subscript t indexes years, in which case months, which are the observation interval of the data, are indicated with $(1/12)$ fractions of a year. The continuously compounded annualized yield on an n-year bond is therefore

$$y_t^{(n)} \equiv -\frac{1}{n} p_t^{(n)}.$$

The natural logarithm of the one-year forward rate at time t for loans between $t+n-1$ and $t+n$ is

$$f_t^{(n)} \equiv p_t^{(n-1)} - p_t^{(n)}.$$

The forward spreads between these forward rates and the one-year yield are

$$fs_t^{(n)} \equiv f_t^{(n)} - y_t^{(1)}.$$

The continuously compounded rate of return from buying an n-year bond at time t and selling it one year later is

$$r_{t+1}^{(n)} \equiv p_{t+1}^{(n-1)} - p_t^{(n)},$$

in which case the excess rate of return is

$$rx_{t+1}^{(n)} \equiv r_{t+1}^{(n)} - y_t^{(1)}.$$

The average of four excess rates of return on bonds with two through five years to maturity is

$$\overline{rx}_{t+1} \equiv (1/4) \sum_{n=2}^{5} rx_{t+1}^{(n)}.$$

Bold symbols without superscripts indicate vectors or matrices. For example, the vector of excess rates of return on bonds with two through five years to maturity is

$$\mathbf{rx}_{t+1} \equiv \left[rx_{t+1}^{(2)},\ rx_{t+1}^{(3)},\ rx_{t+1}^{(4)},\ rx_{t+1}^{(5)}\right]^{\intercal}.$$

When used as right-hand-side variables in a regression, such vectors include a constant. For example,

$$\mathbf{fs}_t \equiv \left[1,\ fs_t^{(2)},\ fs_t^{(3)}, fs_t^{(4)}, fs_t^{(5)}\right]^{\intercal}.$$

Whereas Cochrane and Piazzesi (2005) use the levels of the forward rates as forecasting variables for the excess rates of return on bonds, we follow Cochrane and Piazzesi (2008) and use the averages of the three most recent monthly spreads as the forecasting variables:[1]

$$\overline{\mathbf{fs}}_t \equiv (1/3) \sum_{j=0}^{2} \mathbf{fs}_{t-(j/12)}.$$

The unconstrained forecasting system for the excess rates of return in a particular currency's bond market can therefore be written as

$$\mathbf{rx}_{t+1} = \boldsymbol{\beta}\overline{\mathbf{fs}}_t + \boldsymbol{\varepsilon}_{t+1}, \tag{2.1}$$

[1] Cochrane and Piazzesi (2008) note that levels of forward rates have near unit root components which are unlikely to match up with rational risk premiums. Forward spreads are more likely to be stationary and hence to capture risk premiums. See also the discussion in Cochrane (2015) who advocates using moving averages of forward spreads to avoid spurious predictability due to measurement error in the yields.

where $\boldsymbol{\beta}$ represents the (4×5) matrix of responses of excess returns to the forward spreads. Cochrane and Piazzesi (2005, 2008) motivate their constrained one-factor model of expected bond returns from the finding that the first principal component of the unconstrained expected returns in the system of equations (2.1) explains over 99% of the variance of these expected returns.

This constrained model of a vector of expected returns was first developed by Hansen and Hodrick (1983) and Gibbons and Ferson (1985) who postulated that a set of expected returns could be proportional to a common unobserved factor, v_t:

$$E_t\left(\mathbf{rx}_{t+1}\right) = \mathbf{b}v_t, \tag{2.2}$$

where $\mathbf{b} \equiv [b_2,\ b_3,\ b_4,\ b_5]^\top$. By projecting the unobserved factor onto some observed information, in this case $\overline{\mathbf{fs}}_t$, one can write

$$v_t = \boldsymbol{\gamma}^\top \overline{\mathbf{fs}}_t + \xi_t, \tag{2.3}$$

where by the properties of linear prediction, the error term, ξ_t, is orthogonal to the right-hand-side variables.

Substituting equation (2.3) into equation (2.2) and assuming rational expectations produces a constrained single factor forecasting system that can be written as

$$\mathbf{rx}_{t+1} = \mathbf{b}\boldsymbol{\gamma}^\top \overline{\mathbf{fs}}_t + \boldsymbol{\varepsilon}_{t+1}, \tag{2.4}$$

where $\boldsymbol{\varepsilon}_{t+1}$ now represents both the rational expectations forecast errors for each equation plus $\mathbf{b}\xi_t$. Estimation can be done with the generalized method of moments (GMM) of Hansen (1982) because $\boldsymbol{\varepsilon}_{t+1}$ is orthogonal to $\overline{\mathbf{fs}}_t$ Because $\mathbf{b}$ and $\boldsymbol{\gamma}^\top$ are multiplied together, some identifying constraint must be imposed on the estimation, and we follow Cochrane and Piazzesi (2005) in imposing the constraint on $\mathbf{b}$ that the average of the b_n's equals one:

$$(1/4)\sum_{n=2}^{5} b_n = 1.$$

Whereas the unconstrained model in equation (2.1) has 20 parameters, the constrained model in equation (2.4) has 8 free parameters, 5 in $\boldsymbol{\gamma}$ and 3 in $\mathbf{b}$.

As Cochrane and Piazzesi (2005) note, estimation of the constrained model can be done in two steps. The first step is an OLS regression of the average excess rate of return on the four long-horizon bonds on the average of the forward spreads as in

$$\overline{rx}_{t+1} = \boldsymbol{\gamma}^{\mathsf{T}}\overline{\mathbf{fs}}_t + \bar{\varepsilon}_{t+1}. \tag{2.5}$$

This imposes the constraint that the average of the b_n's equals one.[2] The second step involves OLS regressions without constant terms of three individual excess rates of return on the fitted value from equation (2.5):

$$rx_{t+1}^{(n)} = b_n\left(\widehat{\boldsymbol{\gamma}}^{\mathsf{T}}\overline{\mathbf{fs}}_t\right) + \varepsilon_{t+1}^{(n)}. \tag{2.6}$$

We use the two-year, three-year, and four-year maturities.

2.2.1 The Affine Model with Restrictions

Before discussing the results of estimating the constrained model, we first introduce the affine term structure model that Cochrane and Piazzesi (2008) reverse engineer to be consistent with the forecasting properties from the constrained regressions of excess returns of the long-term bonds on forward spreads.

In a generic ATSM the continuously compounded short-term interest rate is postulated to be a linear function of a K-dimensional vector of state variables, $\mathbf{X}_t$:

$$r_t = \delta_0 + \boldsymbol{\delta}_1^{\mathsf{T}}\mathbf{X}_t.$$

[2] Rather than equal-weighting the excess returns in the first step of the constrained model, Cochrane and Piazzesi (2008) develop an alternative weighting system that relies on an eigenvalue decomposition of the covariance matrix of the expected excess returns from the unrestricted regressions and takes the weights on the excess returns to be the normalized eigenvector associated with the dominant eigenvalue. We follow the approach in the original paper because it is the primary way in which the CP factor has been estimated in the literature, but we note that the alternative approach provides additional insights and may be more useful in understanding other asset markets.

The state variables are assumed to follow a first-order vector autoregression:

$$\mathbf{X}_{t+1} = \boldsymbol{\mu} + \boldsymbol{\Phi}\mathbf{X}_t + \boldsymbol{\Sigma}\boldsymbol{\upsilon}_{t+1}.$$

The vector of innovations, $\boldsymbol{\upsilon}_{t+1}$, is assumed to be $N(0, \boldsymbol{I}_K)$, and the covariance matrix of the state variables is $\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{\mathsf{T}}$. The natural logarithm of the stochastic discount factor is specified to be

$$m_{t+1} = -r_t - \frac{1}{2}\boldsymbol{\lambda}_t^{\mathsf{T}}\boldsymbol{\lambda}_t - \boldsymbol{\lambda}_t^{\mathsf{T}}\boldsymbol{\upsilon}_{t+1}, \tag{2.7}$$

and the innovations to the state variables are thus potential sources of risks. Finally, the prices of these risks are also postulated to be affine functions of the state variables:

$$\boldsymbol{\lambda}_t = \boldsymbol{\lambda}_0 + \boldsymbol{\lambda}_1\mathbf{X}_t,$$

where $\boldsymbol{\lambda}_0$ is $K \times 1$, and $\boldsymbol{\lambda}_1$ is $K \times K$.

The solution of such an affine term structure model uses the basic no-arbitrage asset pricing model,

$$E_t\left(M_{t+1}R_{t+1}^{(n)}\right) = 1, \tag{2.8}$$

where $M_{t+1} = \exp(m_{t+1})$ and $R_{t+1}^{(n)} = \exp\left(r_{t+1}^{(n)}\right)$. Substituting for M_{t+1} and $R_{t+1}^{(n)}$ in equation (2.8) and solving the conditional expectation provides the solution of the ATSM in which the natural logarithms of the bond prices are found to be affine functions of the state variables:

$$p_t^{(n)} = A_n + \boldsymbol{B}_n^{\mathsf{T}}\mathbf{X}_t. \tag{2.9}$$

The recursive formulas for the A_n and $\boldsymbol{B}_n$ coefficients in equation (2.9) are given in Appendix B.

From the solution of the the ATSM, one finds that the expected excess rates of return on bonds are also affine functions of the state variables:

$$E_t\left(rx_{t+1}^{(n)}\right) = -(1/2)\,\boldsymbol{B}_{n-1}^{\top}\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{\top}\boldsymbol{B}_{n-1} + \boldsymbol{B}_{n-1}^{\top}\boldsymbol{\Sigma}\boldsymbol{\lambda}_0 + \boldsymbol{B}_{n-1}^{\top}\boldsymbol{\Sigma}\boldsymbol{\lambda}_1\mathbf{X}_t. \tag{2.10}$$

The three terms on the right-hand side of equation (2.10) are a Jensen's inequality term related to the variance of the rate of return, a constant risk premium, and a time-varying risk premium. In the general ATSM without constraints on the parameters, time-varying expected excess rates of returns on bonds would be driven by the K state variables. This would be inconsistent with the empirical finding that only one state variable is required to forecast economically interesting variation in expected excess returns.

To reconcile the theoretical analysis with the empirical findings, Cochrane and Piazzesi (2008) postulate that the term structure of interest rates depends on four state variables, but they constrain the prices of risks such that only one of these variables drives expected excess rates of return. At least since Litterman and Scheinkman (1991) it has been known that time variation in zero-coupon bond yields can be effectively modeled with the first three principal components of the yields, which are a level effect, l_t, a slope effect, s_t, and a curvature effect, c_t. Hence, these three variables are present as state variables. The fourth state variable is the "return forecasting factor," that is, the CP factor:

$$x_t \equiv \widehat{\boldsymbol{\gamma}}^{\top}\overline{\mathbf{fs}}_t. \tag{2.11}$$

The state vector can therefore be written as $\mathbf{X}_t = (x_t,\ l_t,\ s_t,\ c_t)^{\top}$.[3] Because Cochrane and Piazzesi (2005) empirically find a very strong one-factor structure in the unconstrained model in equation (2.1), Cochrane and Piazzesi (2008) place a set of restrictions on the prices of risks, $\boldsymbol{\lambda}_t$, such that a one-factor structure emerges in equation (2.10). Since x_t is the only variable that can predict expected returns, the columns of the $\boldsymbol{\lambda}_1$ matrix other than the first must be zeros. Cochrane and Piazzesi (2008) also find that the b_n coefficients of the empirical model line up nicely with the

[3]Cieslak and Povala (2015) develop a similar ATSM with three state variables: the expected or 'trend' rate of inflation, a real factor orthogonal to expected inflation, and a forecasting variable that only affects the prices of the two risks.

covariances of the excess returns with the innovations to the level factor. This motivates the full set of restrictions such that

$$\boldsymbol{\lambda}_t = \begin{bmatrix} 0 \\ \lambda_{0l} \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 \\ \lambda_{1l} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_t \\ l_t \\ s_t \\ c_t \end{bmatrix}. \tag{2.12}$$

Thus, although innovations in the four state variables drive the zero-coupon yields and bond prices at all maturities, the only innovation that affects the bond market's stochastic discount factor and hence affects expected rates of return on bonds is the innovation in the level of the term structure, denoted $\upsilon_{l,t+1}$, and the time varying price of this risk is driven by the return forecasting factor. That is,

$$\boldsymbol{\lambda}_t^{\mathsf{T}} \boldsymbol{\upsilon}_{t+1} = \begin{bmatrix} 0 \\ (\lambda_{0l} + \lambda_{1l} x_t)\, \upsilon_{l,t+1} \\ 0 \\ 0 \end{bmatrix}. \tag{2.13}$$

Substituting from equation (2.12) into equation (2.10) gives

$$E_t\left(rx_{t+1}^{(n)}\right) = -(1/2)\, \boldsymbol{B}_{n-1}^{\mathsf{T}} \boldsymbol{\Sigma} \boldsymbol{\Sigma}^{\mathsf{T}} \boldsymbol{B}_{n-1} + \boldsymbol{B}_{n-1}^{\mathsf{T}} \boldsymbol{\Sigma} \begin{bmatrix} 0 \\ (\lambda_{0l} + \lambda_{1l} x_t) \\ 0 \\ 0 \end{bmatrix}. \tag{2.14}$$

While equation (2.14) is quite close to the constrained econometric model in equation (2.4) in that each expected return loads with a different coefficient onto the common forecasting factor, the constrained model makes the additional assumption that the constant terms in the equations share the same proportionality as the slope coefficients. The Jensen's inequality terms do not scale in the same way, which makes this assumption technically incorrect. Because these terms are generally considered to be small, in what follows we ignore this issue and follow the approach of Cochrane and Piazzesi (2008).[4]

2.3 Estimation Results for Nine Term Structures

In this section we estimate the Cochrane and Piazzesi (2005) empirical model for the zero-coupon government bond yields of nine of the G-10 currencies: the Australian dollar, AUD; the Canadian dollar, CAD; the Swiss franc, CHF; the euro, EUR, spliced with data from the Deutsche mark; the British pound, GBP; the Japanese yen, JPY; the Norwegian krone, NOK; the Swedish krona, SEK; and the USD. After reviewing the available term structure data for the New Zealand dollar, we viewed it as unreliable and therefore did not include it in our analysis. Sources of data are described in Appendix A.

We present the results in two sections corresponding to samples of data that would have been available when the original model was first estimated and samples of data that subsequently became available. Because the last observation on the dependent variable in the the first data set is December 2003, we refer to these data as the pre-2004 sample. We begin observations on the dependent variable in the second data set in December 2004 to avoid overlap with the first data set, and we refer to these data as the post-2003 sample. To allow for samples that coincide with the exchange rate data, the dependent variables for the first sample begin in 1974:12 for the USD, the GBP, and the EUR; in 1989:03 for the CHF; in 1987:03 for the CAD; in 1986:03 for the JPY; in 1988:04 for the AUD; in 1988:03 for the SEK; and in 1999:03 for the NOK. The first sample is particularly short for the NOK, so we do not think those results are particularly informative, but we choose to include the results simply because the NOK is included in the post-2003 analysis.

[4] The Online Appendix presents some results of a model that relaxes this restrictive assumption by allowing for separate constants at each maturity. We find that the relevant parameters associated with the time-varying forecasts of the two models are quite close and inference is quite similar.

2.3.1 Results with Pre-2004 Data

Table 2.1 reports the estimation of the constrained model in equation (2.4) with the two-step OLS procedure described above.

Both samples are included in Table 2.1 with the notation CUR1 or CUR2 indicating the currency of denomination of the bonds and either the first or second sample period, respectively. We report asymptotic GMM standard errors that account for the overlapping forecasts and the fact that the second step in the estimation uses estimated coefficients from the first step.[5]

Although the unconstrained results are not reported because of the large number of parameters, the first thing to notice in Table 2.1 is the strong support for the single factor forecasting structure of expected excess returns in each of the nine term structures in the unconstrained estimations. The far right column labeled $\%PC1$ presents the proportion of the variance of the four unconstrained estimates of the excess rates of return, denominated in the particular currency of that row, that is explained by their first principal component. For all the currencies in the first sample, the first principal component explains at least 98.8% of the variance of these expected excess returns. This evidence represents strong support for the one-factor forecasting model of expected excess bond returns in each of the currencies.

The second noteworthy aspect of Table 2.1 is the remarkable similarity in the coefficient estimates of b_2, b_3, and b_4 across currencies. The estimated values of b_2 range from 0.37 for the JPY to 0.47 for the CHF. The estimated values of b_3 range from 0.80 for the SEK to 0.87 for the AUD. The estimated values of b_4 range from 1.19 for the CHF to 1.23 for the USD and the JPY. From equation (2.14) we see that the estimated values of the b_n's in an ATSM differ because of the different values of $\mathbf{B}_{\mathbf{n-1}}^{\intercal}\boldsymbol{\Sigma}$ associated with the CP factor. The recursive solution for the $\mathbf{B}_{\mathbf{n}}^{\intercal}$ in equation (2.32) indicates that values of $\mathbf{B}_{\mathbf{n}}^{\intercal}$ change as $\boldsymbol{\Phi}^{*}$, the risk neutral autocorrelation matrix of the state variables, is raised to higher powers. Thus, the finding of similar values of the b_n's

[5] The standard errors could be constructed as in Hansen and Hodrick (1980), by equally weighting the 11 lagged covariances that are non-zero by construction when forecasting annual excess returns with overlapping monthly data. These standard errors are not guaranteed to be positive definite, and in fact in some cases they were not. Consequently, we rely on Newey and West (1987) standard errors using 18 lags as in Cochrane and Piazzesi (2005).

Table 2.1: The Single Factor Cochrane and Piazzesi (2005) Model

CUR	b_2	b_3	b_4	γ_1	γ_2	γ_3	γ_4	γ_5	$\chi^2(4)$	R^2	%PC1
USD1	0.46	0.86	1.23	-0.04	2.32	3.65	1.73	-4.67	100.10	0.39	0.997
	(0.02)	(0.02)	(0.01)	(0.63)	(1.63)	(1.33)	(0.95)	(0.58)	⟨0.00⟩		
USD2	0.35	0.78	1.24	0.06	-7.57	2.17	2.57	-0.80	25.44	0.37	0.973
	(0.06)	(0.06)	(0.01)	(0.69)	(2.61)	(1.76)	(0.60)	(0.61)	⟨0.00⟩		
GBP1	0.44	0.86	1.21	1.04	-5.40	11.24	-6.54	0.39	9.46	0.09	0.996
	(0.05)	(0.03)	(0.02)	(0.53)	(4.31)	(11.72)	(14.23)	(5.97)	⟨0.05⟩		
GBP2	0.37	0.82	1.23	0.17	-8.25	-5.32	27.01	-15.88	67.63	0.29	0.964
	(0.08)	(0.06)	(0.02)	(0.76)	(7.48)	(19.84)	(24.71)	(10.43)	⟨0.00⟩		
EUR1	0.44	0.86	1.21	1.02	-1.77	-4.19	16.47	-10.50	9.66	0.12	0.992
	(0.03)	(0.03)	(0.01)	(0.90)	(5.07)	(14.33)	(19.38)	(9.17)	⟨0.05⟩		
EUR2	0.34	0.77	1.22	1.61	-21.01	35.23	-24.91	6.39	6.56	0.08	0.973
	(0.06)	(0.05)	(0.01)	(1.49)	(14.67)	(33.29)	(35.70)	(14.36)	⟨0.16⟩		
CHF1	0.47	0.86	1.19	0.78	-1.45	-9.62	26.03	-14.60	22.24	0.24	0.994
	(0.08)	(0.04)	(0.03)	(1.45)	(11.88)	(30.09)	(38.64)	(18.01)	⟨0.00⟩		
CHF2	0.28	0.74	1.25	0.12	-10.41	16.56	-14.21	5.73	164.11	0.40	0.983
	(0.06)	(0.06)	(0.01)	(0.47)	(2.88)	(8.60)	(11.62)	(5.46)	⟨0.00⟩		
CAD1	0.46	0.86	1.20	1.64	0.53	-5.16	19.68	-13.59	34.29	0.23	0.993
	(0.03)	(0.03)	(0.01)	(0.84)	(2.34)	(5.25)	(6.95)	(4.03)	⟨0.00⟩		
CAD2	0.38	0.81	1.22	1.59	3.77	-13.41	21.50	-10.71	17.68	0.16	0.948
	(0.07)	(0.06)	(0.02)	(1.04)	(3.60)	(6.40)	(6.49)	(2.78)	⟨0.00⟩		
JPY1	0.37	0.81	1.23	1.00	-14.98	17.33	-0.29	-5.00	25.87	0.25	0.997
	(0.02)	(0.02)	(0.00)	(1.23)	(3.76)	(4.16)	(2.06)	(1.87)	⟨0.00⟩		
JPY2	0.30	0.73	1.29	0.40	3.33	-1.04	4.04	-3.31	36.36	0.34	0.994
	(0.03)	(0.04)	(0.02)	(0.22)	(1.75)	(2.70)	(1.74)	(1.12)	⟨0.00⟩		
AUD1	0.44	0.87	1.21	0.82	-32.92	107.79	-138.85	61.21	15.94	0.11	0.996
	(0.02)	(0.02)	(0.01)	(1.07)	(12.15)	(42.03)	(61.52)	(29.98)	⟨0.00⟩		
AUD2	0.33	0.81	1.22	1.50	3.03	2.68	-4.15	0.90	17.56	0.12	0.970
	(0.11)	(0.08)	(0.02)	(1.02)	(2.59)	(2.08)	(1.46)	(1.11)	⟨0.00⟩		

(*Continued*)

CUR	b_2	b_3	b_4	γ_1	γ_2	γ_3	γ_4	γ_5	$\chi^2(4)$	R^2	%PC1
SEK1	0.39	0.80	1.21	1.27	30.34	-68.63	73.43	-29.16	5.51	0.11	0.998
	(0.05)	(0.03)	(0.02)	(1.43)	(13.92)	(39.50)	(51.60)	(23.65)	⟨0.24⟩		
SEK2	0.37	0.80	1.22	1.87	15.10	-40.53	47.55	-19.82	11.40	0.12	0.974
	(0.08)	(0.06)	(0.02)	(1.10)	(8.60)	(17.47)	(19.23)	(8.42)	⟨0.02⟩		
NOK1	0.44	0.87	1.21	1.97	32.85	-127.27	200.04	-101.71	116.38	0.42	0.987
	(0.03)	(0.02)	(0.01)	(1.23)	(19.32)	(69.50)	(98.33)	(46.53)	⟨0.00⟩		
NOK2	0.37	0.79	1.22	0.99	-7.11	6.46	-2.84	0.63	2.65	0.07	0.870
	(0.13)	(0.12)	(0.02)	(0.97)	(10.76)	(23.26)	(23.36)	(8.99)	⟨0.62⟩		

Description: The Table reports coefficient estimates for the two-step estimation of the constrained single factor model for two sample periods. The sample periods for the dependent variables in the first sample labeled CUR1 all end in 2003:12. These samples begin in 1974:12 for the USD, the GBP, and the EUR; in 1989:03 for the CHF; in 1987:03 for the CAD; in 1986:03 for the JPY; in 1988:04 for the AUD; in 1988:03 for the SEK; and in 1999:03 for the NOK. For the second sample period labeled CUR2 the dependent variables for all currencies begin in 2004:12 and end in 2016:12. The first step in the estimation involves OLS regressions of the average one-year excess rates of return on bonds with two through five years to maturity, $\overline{rx}_{t+1}$, on a constant and the average of the current value and two monthly lags of the four forward spreads, $\overline{\mathbf{fs}}_t$:

$$\overline{rx}_{t+1} = \gamma^{\intercal}\overline{\mathbf{fs}}_t + \overline{\varepsilon}_{t+1}.$$

The second step involves OLS regressions of the individual excess rates of return on bonds with two through four years to maturity on the fitted value from the first step:

$$rx^{(n)}_{t+1} = b_n\left(\widehat{\gamma}^{\intercal}\overline{\mathbf{fs}}_t\right) + \varepsilon^{(n)}_{t+1}$$

Standard errors in the first step are based on the usual GMM versions from OLS orthogonality conditions, and the standard errors in the second step allow for the estimation error in the first step. All standard errors are constructed with 18 Newey-West (1987) lags and are in parentheses. The $\chi^2(4)$ statistic tests the hypothesis that γ_2 through γ_5 equal zero with p-values in angled brackets. The R^2 is from the first step regression. The column labeled $\%PC1$ presents the percentage of the variance of the unconstrained estimates of the four excess rates of return explained by their first principal component.

Interpretation: There is strong evidence for a single factor forecasting structure of expected excess returns for all currencies in each of the two non-overlapping samples. The coefficient estimates of the **b**'s are remarkably similar across currencies and across the two samples. The estimates of the $\boldsymbol{\gamma}$'s are not similar across currencies, and we observe large differences in the $\boldsymbol{\gamma}$ estimates across the two samples.

across countries indicates that if we were to estimate an ATSM for each currency, the resulting $\boldsymbol{\Phi}^*$ estimates would be quite similar across currencies. At this point, we leave this as a conjecture for future research.

While there is considerable variety in the estimates of the γ_j's across the different currencies, the $\chi^2(4)$ statistics for all currencies except the SEK provide strong rejections of the currency-by-currency null hypothesis that the time-varying, right-hand-side variables have no collective forecasting power. Particularly large values of coefficients for the AUD, SEK, and NOK are an indication of multicollinearity. Although Cochrane and Piazzesi (2005) found a clear "tent" pattern in their projection of average returns onto the levels of the five forward rates, we only see this pattern in projections onto the four forward spreads for the USD and JPY data.

There are at least three reasons why the estimates of the $\boldsymbol{\gamma}$'s might differ across currencies. The first explanation takes a rational expectations econometrics view and recognizes that the forward spreads capture the risk exposures of a country as represented by the reduced form coefficients from an ATSM. Underlying structural differences in the nature of risks would consequently manifest themselves in different $\boldsymbol{\gamma}$'s. Monetary and fiscal policies certainly differ across countries, and we do not attempt to relate the underlying coefficients of the ATSM to more structural coefficients in equations such as the Taylor (1993) rule.

Alternatively, a second reason would take the perspective of Bekaert, Hodrick, and Marshall (2001) who argue that the rational expectations econometrics perspective may be too strong. Developed countries, such as those studied here, may actually be following the same time series rule, but the realizations of the shocks hitting the economies may have differed across countries. It may take a very long sample for a particular economy to experience all of the possible realizations from the policy rule with their ex ante frequencies that investors anticipated during the sample. It is certainly true that ex post experiences with inflation have differed across the countries, although at a casual level, all countries now seem to be converging to relatively low rates of expected inflation.

As an example of this last perspective, it is notable in Table 2.1 that the R^2's from the first-step regression of the average return on the forward spreads are the highest for the USD and

JPY. Bekaert, Hodrick, and Marshall (2001) argue that the decline in U.S. inflation under Federal Reserve Chairmen Volcker and Greenspan represents a one-sided realization that made the ex post returns on investments in long-term bonds better than was anticipated.[6] Inflation in Japan during much of the sample was also surprisingly low. Thus, the Japanese situation could be similar to the U.S. in that the stagnation in the Japanese economy and its ultimate experiences with deflation resulted in surprisingly good ex post returns on long-term Japanese bonds even though bond yields were quite low to start.

A third reason is offered by Cochrane and Piazzesi (2008) who note that the construction of the CP factor is sensitive to how the term structure data are derived. While the USD data are constructed using actual prices and the bootstrap method of Fama and Bliss (1987), zero-coupon term structure data on other currencies result from sequential, cross-sectional estimation using the flexible functional form approaches of Nelson and Siegel (1987) or Svensson (1995). In comparing forecasts of USD data from Gürkaynak, Sack, and Wright (2007) that are constructed from the flexible form of Svensson (1995) to forecasts from the Fama and Bliss (1987) data, Cochrane and Piazzesi (2008) find evidence of multicollinearity in the former and more predictive power in the latter. Consistent with these findings, we noted above that several of the currencies show evidence of multicollinearity in the forecasting variables.

2.3.2 Results with Post-2003 Data

We now discuss the results for the second sample from 2004 to 2016 that are also presented in Table 2.1 in the rows labeled CUR2. While the one-factor structure of expected excess returns, estimated from unconstrained regressions, is not quite as strong in this sample, we still see that the first principal component of the unconstrained expected returns explains between 86.6% of the variance for the NOK and 99.4% for the JPY. The similarity in the coefficient estimates of b_2, b_3, and b_4 is maintained, and differences from the estimates in the first sample are generally small. The estimated values of b_2 now range from 0.28 for the CHF to 0.38 for the CAD; the estimated

[6] See Bauer and Rudebusch (2017) for an analysis of the U.S. term structure that allows for declining stochastic trends in both the long-run expected rate of inflation and the equilibrium real interest rate.

values of b_3 range from 0.73 for the JPY to 0.82 for the GBP and the CAD; and the estimated values of b_4 range from 1.22 for the EUR, CAD, AUD, SEK, and NOK to 1.29 for the JPY. The estimates of the $\boldsymbol{\gamma}$'s appear much different in the second sample compared to the first. Differences are particularly large for the GBP, EUR, CHF, JPY, AUD, and NOK.

As a first step in analyzing the out-of-sample performance of the Cochrane and Piazzesi (2005) model, Table 2.2 presents tests of the equality of the **b**'s and $\boldsymbol{\gamma}$'s across the two samples on a currency-by-currency basis.

For the **b**'s, even though the coefficient estimates are quite similar across the two samples, their small standard errors lead to rejections of equality of the three coefficients for the EUR at the 1% marginal level of significance, for the CHF at the 3% level, and for the JPY at smaller than the 1% level. The tests of the $\boldsymbol{\gamma}$'s reject equality of the parameter estimates across the two periods for the USD, the JPY, and the NOK at less than the 1% level, for the GBP at the 9% level, and for the AUD at the 10% level. These findings provide the first evidence of instability in the forecasting relations.

2.3.3 Correlation Matrix and Variance Decomposition of Country CP Factors

Because one-factor forecasting structures characterize each of the term structures quite well, a natural question to ask is how correlated are the various CP factors. Table 2.3 provides correlation matrices for the respective currency-specific CP factors for the two sample periods.

Panel A of Table 2.3 presents the results for the first sample, and we find that 26 of the 36 correlations are positive, but only the GBP-CHF correlation of .63 is larger than .5. Of the nine negative ones, the JPY-NOK correlation is the most negative at -.30. The last column in Table 2.3 labeled $\%PC(i)$ reports the percent of the variance of the nine country-specific CP factors that is explained by the respective principal components. The first three principal components explain 82% of the total variance. While this evidence is suggestive that global risk factors may be at work in explaining the ability of the CP factors to forecast excess bond returns, it is certainly not

Table 2.2: Tests of Equality of Coefficients for the Two Samples

CUR	$\chi^2(3)$ for **b**'s	$\chi^2(5)$ for $\boldsymbol{\gamma}$'s
USD	5.15	81.77
	⟨0.16⟩	⟨0.00⟩
GBP	4.77	9.52
	⟨0.19⟩	⟨0.09⟩
EUR	11.64	4.39
	⟨0.01⟩	⟨0.49⟩
CHF	8.69	5.47
	⟨0.03⟩	⟨0.36⟩
CAD	5.46	7.16
	⟨0.14⟩	⟨0.21⟩
JPY	26.52	25.87
	⟨0.00⟩	⟨0.00⟩
AUD	2.83	9.36
	⟨0.42⟩	⟨0.10⟩
SEK	1.12	1.27
	⟨0.77⟩	⟨0.94⟩
NOK	5.84	35.65
	⟨0.12⟩	⟨0.00⟩

Description: The Table reports two statistics that test the equality of the coefficient estimates in the Cochrane and Piazzesi (2005) models for the two samples estimated in Table 2.1. The first test examines the **b** coefficients and is distributed as a $\chi^2(3)$. The second test examines the $\boldsymbol{\gamma}$ coefficients and is distributed as a $\chi^2(5)$. The p-values are in angled brackets. The first sample periods for the dependent variables all end in 2003:12. These samples begin in 1974:12 for the USD, the GBP, and the EUR; in 1989:03 for the CHF; in 1987:03 for the CAD; in 1986:03 for the JPY; in 1988:04 for the AUD; in 1988:03 for the SEK; and in 1999:03 for the NOK. The second sample period for the dependent variables is 2004:12 to 2016:12 for all currencies.

Interpretation: These findings provide the first evidence of instability in the forecasting relations. While we cannot reject the hypotheses that the **b**'s and $\boldsymbol{\gamma}$'s are equal across the two samples for the CAD and the SEK, for the other currencies we find strong rejections of equality of the **b**'s for the EUR, the CHF, and the JPY; and we find strong rejections of equality of the $\boldsymbol{\gamma}$'s for the USD, the JPY, and the NOK, and slightly weaker evidence of inequality of the $\boldsymbol{\gamma}$'s for the GBP and the AUD.

Table 2.3: Correlation Matrices and Variance Decompositions of Country CP Factors

Panel A: Pre-2004 Data										
	GBP	EUR	CHF	CAD	JPY	AUD	SEK	NOK	i	%PC(i)
USD	0.06	0.23	0.26	0.37	0.22	0.35	0.13	-0.02	1	0.40
GBP	1.00	-0.19	0.63	0.49	0.45	0.10	-0.24	0.27	2	0.29
EUR		1.00	0.12	-0.11	0.18	-0.08	0.44	0.44	3	0.13
CHF			1.00	0.19	0.34	-0.16	-0.20	0.34	4	0.09
CAD				1.00	0.24	0.12	0.10	-0.07	5	0.04
JPY					1.00	0.02	0.03	-0.30	6	0.03
AUD						1.00	-0.09	0.18	7	0.01
SEK							1.00	0.37	8	0.01
NOK								1.00	9	0.00
Panel B: Post-2003 Data										
	GBP	EUR	CHF	CAD	JPY	AUD	SEK	NOK	i	%PC(i)
USD	0.42	-0.07	0.12	0.34	0.35	0.12	0.12	0.18	1	0.37
GBP	1.00	0.39	0.34	0.40	-0.11	0.08	0.19	0.49	2	0.18
EUR		1.00	0.33	0.19	-0.32	0.08	0.10	0.35	3	0.14
CHF			1.00	0.17	0.08	-0.01	0.14	0.44	4	0.09
CAD				1.00	0.37	0.15	0.12	0.03	5	0.08
JPY					1.00	0.06	-0.00	-0.29	6	0.06
AUD						1.00	0.14	-0.32	7	0.04
SEK							1.00	0.13	8	0.02
NOK								1.00	9	0.01

Description: The Table presents the correlation matrices of the CP factors, the fitted return forecasting variables from the term structure regressions for the different currencies and the two sample periods in Table 2.1. In Panel A, the sample periods for the dependent variables all end in 2003:12. The samples begin in 1974:12 for the USD, the GBP, and the EUR; in 1989:03 for the CHF; in 1987:03 for the CAD; in 1986:03 for the JPY; in 1988:04 for the AUD; in 1988:03 for the SEK; and in 1999:03 for the NOK. Because the samples for the different currencies are of different lengths, the correlations are estimated over the shorter of the two samples. In Panel B, the sample periods for the dependent variables are all 2004:12 to 2016:12. The last column labeled %PC reports the percent of variance explained by the *i*-th principal component.

Interpretation: Although the estimated correlations between the different currency CP factors are generally positive in both samples, they are relatively small in magnitude, and in the post-2003 data, twelve out of 36 correlations change sign from the first sample. In the first sample, the first three principal components explain 82% of the correlation matrix, which suggests that global risk factors may explain the ability of the CP factors to forecast excess bond returns, but such evidence is certainly not definitive. In the second sample, the share of variance explained by the first three principal components drops to 69%, suggesting instability in the model.

definitive.[7]

When we examine the post-2003 samples in Panel B of Table 2.3, we find that six of the 36 correlations are negative, and the largest positive correlation is now the GBP-NOK correlation of .54, which is the only correlation greater than .5. Twelve of the correlations change sign, and the largest switch is the GBP-EUR correlation which increased from -.19 to .39. The share of the variance explained by the first three principal components falls to 69%. These changes in correlations are another indication of instability in the model.

We will examine out-of-sample forecasting of bond returns below, but first, we examine some international implications of the model.

2.4 International Implications

This section derives some implications of the Cochrane and Piazzesi (2005, 2008) model for foreign exchange markets. Doing so requires the introduction of subscripts for the currencies, and we subscript the USD variables with a one and variables denominated in an arbitrary foreign currency with a j. We define exchange rates as $S_{ij,t}$, which represents the currency j price of base currency i at time t. The continuously compounded rate of appreciation of base currency i relative to currency j between times t and $t+1$ is denoted $\Delta s_{ij,t+1}$.

We first argue that tight restrictions between the term structure models of the two currency markets and the relative rate of currency appreciation are not supported empirically.[8] Then, we consider some less constrained empirical predictions.

To understand this argument, consider the basic no arbitrage asset pricing equation for a particular currency i that prices all returns denominated in that currency as in equation (2.8); but now, let $Q_{i,t+1}$ represent this general SDF that prices these generic returns, $R_{i,t+1}$, which include returns on

[7] Jotikasthira, Le, and Lundblad (2015) investigate the determinants of the correlations across several major currency term structures.

[8] See Backus, Foresi, and Telmer (2001) for a discussion of the links between fully specified SDF's and the rate of currency appreciation when financial markets are complete, and see Brandt and Santa-Clara (2002) and Brandt, Cochrane, and Santa-Clara (2006) for discussions of the effects of incomplete markets.

all currency i denominated assets and not just the bond returns of equation (2.8). Thus, we have

$$E_t\left(Q_{i,t+1}R_{i,t+1}\right) = 1. \tag{2.15}$$

The difference between the SDF in equation (2.15), $Q_{i,t+1}$, that prices all currency i returns and the SDF in equation (2.8), $M_{i,t+1}$, that only prices the currency i bond returns is that $Q_{i,t+1}$ can contain risks that are orthogonal to the risks that are priced in the term structure of interest rates. Analytically, we can decompose $Q_{i,t+1}$ as

$$Q_{i,t+1} = M_{i,t+1}Z_{i,t+1}. \tag{2.16}$$

Consistency of the two no arbitrage conditions requires that $E_t\left(Z_{i,t+1}\right) = 1$, because the risk free rate is correctly priced by $M_{i,t+1}$; $E_t\left(M_{i,t+1}Z_{i,t+1}\right) = E_t\left(M_{i,t+1}\right)E_t\left(Z_{i,t+1}\right)$, because $Z_{i,t+1}$ and $M_{i,t+1}$ are orthogonal; and for the return on an n-period bond, $R^{(n)}_{i,t+1}$, $E_t\left(Z_{i,t+1}R^{(n)}_{i,t+1}\right) = E_t\left(Z_{i,t+1}\right)E_t\left(R^{(n)}_{i,t+1}\right)$ because $M_{i,t+1}$ contains all risks priced in the bond market making $Z_{i,t+1}$ orthogonal to $R^{(n)}_{i,t+1}$.

2.4.1 Implications for the Innovation in Currency Appreciation

If financial markets are complete, it is well known that there is a tight relation between the rate of appreciation of currency i relative to currency j and the difference between the natural logarithms of the stochastic discount factor of currency i, $q_{i,t+1}$, and the stochastic discount factor of currency j, $q_{j,t+1}$:

$$\Delta s_{ij,t+1} = q_{i,t+1} - q_{j,t+1}. \tag{2.17}$$

From equation (2.17) we see that the innovation in the rate of appreciation of currency i relative to currency j should be completely explained by the difference in the innovations in the risks present in the natural logarithms of the two currencies stochastic discount factors.

To draw out the international implications of the Cochrane and Piazzesi (2008) ATSM, we substitute for the q's in equation (2.17) to find

$$\Delta s_{ij,t+1} = m_{i,t+1} + z_{i,t+1} - m_{j,t+1} - z_{j,t+1}, \tag{2.18}$$

where $z_{i,t} \equiv \ln(Z_{i,t})$. Notice that if the Cochrane and Piazzesi (2008) ATSM correctly characterized the term structure in each currency, if asset markets were complete, and if the term structure SDF's contained all the sources of risks, then the z's could be eliminated from equation (2.18). After substituting for the innovations in the m's from equation (2.13), the innovation in the rate of appreciation of currency i relative to currency j would be

$$\Delta s_{ij,t+1} - E_t\left(\Delta s_{ij,t+1}\right) = \left(\lambda_{j,0l} + \lambda_{j,0l} x_{j,t}\right) \upsilon_{j,l,t+1} - \left(\lambda_{i,0l} + \lambda_{i,0l} x_{i,t}\right) \upsilon_{i,l,t+1}. \tag{2.19}$$

Thus, the innovation in $\Delta s_{ij,t+1}$ would be fully explained by the innovations in $m_{j,t+1}$ and $m_{i,t+1}$. In the Cochrane and Piazzesi (2008) ATSM, the innovation in the SDF of a currency is the innovation in the level factor of the term structure interacted with a constant and with the predetermined CP factor.

We investigate this issue for rates of appreciation of the USD versus the other eight currencies in Table 2.4. Because the exact fit of equation (2.19) would be unlikely to hold, we run regressions with the expectation that if the model were true, we would have quite significant explanatory power. We proxy the innovation in the rate of appreciation of the USD with respect to currency j with the excess rate of return on a USD investment in the currency j money market, $-\Delta s_{1j,t+1} + r_{j,t} - r_{1,t}$. We proxy the innovations in the level factors with the changes in the levels, as represented by the first principal components of the term structures, because these first principal components are highly serially correlated. For simplicity, we also just report results for the full sample periods associated with each currency.

In the regressions in Table 2.4 the R^2's range from 2% for the CAD and the CHF to 23% for the JPY. This represents strong evidence that the constrained Cochrane and Piazzesi (2008) term structure models do not span the spaces of risks that characterize the rates of currency depreciation,

Table 2.4: Explaining Currency Market Excess Returns with Changes in Level Factors

CUR j	β_0	β_1	β_2	β_3	β_4	R^2
GBP	0.67	0.56	-0.13	-0.64	0.48	0.04
	(1.49)	(0.64)	(0.22)	(0.58)	(0.22)	
EUR	-0.37	-0.70	-0.02	0.63	-0.47	0.04
	(1.65)	(0.59)	(0.25)	(1.11)	(0.57)	
CHF	0.04	-0.10	-0.34	0.03	-0.06	0.02
	(1.89)	(1.06)	(0.41)	(0.85)	(0.65)	
CAD	1.58	0.27	-0.02	0.17	0.15	0.02
	(1.36)	(0.87)	(0.30)	(0.77)	(0.22)	
JPY	-0.97	1.51	-0.78	-3.02	0.90	0.22
	(2.05)	(1.03)	(0.23)	(1.26)	(0.48)	
AUD	3.91	0.09	-0.28	0.84	0.49	0.13
	(2.05)	(0.92)	(0.39)	(1.49)	(0.65)	
SEK	0.65	0.31	-0.68	2.35	-0.47	0.20
	(1.78)	(1.32)	(0.61)	(0.74)	(0.49)	
NOK	1.01	-0.04	-1.17	1.26	-0.07	0.12
	(2.11)	(1.55)	(0.71)	(3.31)	(4.05)	

Description: The Table presents estimation results for the regression,

$$-\Delta s_{1j,t+1} + r_{j,t} - r_{1,t} = \beta_0 + \beta_1 \Delta l_{1,t+1} + \beta_2 \Delta l_{1,t+1} \cdot x_{1,t} + \beta_3 \Delta l_{j,t+1} + \beta_4 \Delta l_{j,t+1} \cdot x_{j,t} + \varepsilon_{1j,t+1},$$

where the dependent variable is the excess rate of return in USD on an annual investment in the money market of currency j, which is our proxy for the innovation in the rate of dollar appreciation. The regressors are the contemporaneous changes in the first principal components of the yields for the USD, $\Delta l_{1,t+1}$, and for currency j, $\Delta l_{j,t+1}$, and the interaction of these variables with their respective currency-specific CP factors, which are the term structure excess return forecasting variables, $x_{1,t}$ and $x_{j,t}$. Standard errors are in parentheses. The sample periods for the dependent variables all end in 2016:12. The samples begin in 1974:12 for the GBP and the EUR; in 1989:03 for the CHF; in 1987:03 for the CAD; in 1986:03 for the JPY; in 1988:04 for the AUD; in 1988:03 for the SEK; and in 1999:03 for the NOK. The samples all end in 2016:12.

Interpretation: The small R^2's in regressions for models proxying the predictions of the Cochrane and Piazzesi (2008) ATSM suggest that these variables do not span the spaces of risks that characterize the rates of currency depreciation, suggesting the possible presence of additional risks that price all assets.

which we interpret as evidence for the presence of additional risks in the SDF's that price all assets.

These results are consistent with the analysis of Sarno, Schneider, and Wagner (2012) who estimate four-factor, latent variable ATSM's for the bond markets of two currencies and find that while the bonds are priced very well, the variation of the rate of currency appreciation from the implied ATSM stochastic discount factors does not match well with the actual rate of currency appreciation. Of course, the results could also indicate that financial markets are incomplete as in the analysis of Brandt and Santa-Clara (2002), but the exchange rate volatility puzzle first discussed in Brandt, Cochrane, and Santa-Clara (2006) suggests that the stochastic discount factors should be relatively highly correlated.

2.4.2 Implications for expected cross-currency investments

To investigate expected rates of return on cross-currency investments that are implied by the Cochrane and Piazzesi (2008) ATSM model but with $Z_{i,t+1}$ present, let $Z_{i,t+1}$ be log-normally distributed. Then, we can assume that the stochastic process for $z_{i,t+1}$ is

$$z_{i,t+1} = -\frac{1}{2}\boldsymbol{\lambda}_{z_i,t}^{\intercal}\boldsymbol{\lambda}_{z_i,t} - \boldsymbol{\lambda}_{z_i,t}^{\intercal}\boldsymbol{\upsilon}_{z_i,t+1}, \tag{2.20}$$

where $\boldsymbol{\upsilon}_{z_i,t+1}$ is a vector of risks that are distributed $N(0, I)$ and that are orthogonal to the vector of risks, $\boldsymbol{\upsilon}_{i,t+1}$, that drive the term structure of interest rates in that currency.

Substituting for the SDF's in equation (2.18) from equations (2.7) and (2.20) and rearranging terms gives the excess rate of return in currency i on a one-year investment in the money market of currency j:

$$\begin{aligned} -\Delta s_{ij,t+1} + r_{j,t} - r_{i,t} = {} & \frac{1}{2}\left(\boldsymbol{\lambda}_{i,t}^{\intercal}\boldsymbol{\lambda}_{i,t} - \boldsymbol{\lambda}_{j,t}^{\intercal}\boldsymbol{\lambda}_{j,t}\right) + \frac{1}{2}\left(\boldsymbol{\lambda}_{z_i,t}^{\intercal}\boldsymbol{\lambda}_{z_i,t} - \boldsymbol{\lambda}_{z_j,t}^{\intercal}\boldsymbol{\lambda}_{z_j,t}\right) \\ & + \\ & \boldsymbol{\lambda}_{i,t}^{\intercal}\boldsymbol{\upsilon}_{i,t+1} - \boldsymbol{\lambda}_{j,t}^{\intercal}\boldsymbol{\upsilon}_{j,t+1} + \boldsymbol{\lambda}_{z_i,t}^{\intercal}\boldsymbol{\upsilon}_{z_i,t+1} - \boldsymbol{\lambda}_{z_j,t}^{\intercal}\boldsymbol{\upsilon}_{z_j,t+1}. \end{aligned} \tag{2.21}$$

Taking the conditional expectation of equation (2.21) gives

$$\mathrm{E}_t\left(-\Delta s_{ij,t+1} + r_{j,t} - r_{i,t}\right) = \frac{1}{2}\left(\boldsymbol{\lambda}_{i,t}^{\top}\boldsymbol{\lambda}_{i,t} - \boldsymbol{\lambda}_{j,t}^{\top}\boldsymbol{\lambda}_{j,t}\right) + \frac{1}{2}\left(\boldsymbol{\lambda}_{z_i,t}^{\top}\boldsymbol{\lambda}_{z_i,t} - \boldsymbol{\lambda}_{z_j,t}^{\top}\boldsymbol{\lambda}_{z_j,t}\right). \tag{2.22}$$

The right-hand side of equation (2.22) is the expected excess rate of return to borrowing one unit of currency i, investing that amount in the currency j money market, and bearing the foreign exchange risk.[9] By imposing the constraints of the one-factor forecasting model for the two bond markets in equation (2.12), we find

$$\boldsymbol{\lambda}_{j,t}^{\top}\boldsymbol{\lambda}_{j,t} = \left(\lambda_{j,0l} + \lambda_{j,1l}x_{j,t}\right)^2 = \lambda_{j,0l}^2 + 2\lambda_{j,0l}\lambda_{j,1l}x_{j,t} + \lambda_{j,1l}^2 x_{j,t}^2. \tag{2.23}$$

Substituting from equation (2.23), for both currencies i and j, into equation (2.22) implies that the CP forecasting factors, $x_{i,t}$ and $x_{j,t}$, from the bond markets of the two currencies and their squared values should forecast the excess rate of return to investing a unit of currency i in the currency j money market while bearing the foreign exchange risk.

In contrast to the predictability implied by this ATSM approach, the literature on uncovered interest rate parity predicts that the excess rate of return to investing a unit of currency i in the currency j money market should be unpredictable, and empirical evidence that rejects this hypothesis often focuses on the interest differential, $r_{j,t} - r_{i,t}$, as a predictor, as in the analysis of Fama (1984).

To investigate the predictability of excess currency returns we blend these two specifications as in the following forecasting regression in which we use the USD as the base currency:

$$-\Delta s_{1j,t+1} + r_{j,t} - r_{1,t} = \phi_0 + \phi_1 x_{1,t} + \phi_2 x_{j,t} + \phi_3(r_{j,t} - r_{1,t}) + \epsilon_{1j,t+1}^{s}. \tag{2.24}$$

In equation (2.24) we leave out the squared values of the CP forecasting factors as these are numerically small, and we also do not impose any constraints on the regression coefficients of the CP factors that arise from the term structure theory because we do not observe $\boldsymbol{\lambda}_{z_i,t}^{\top}\boldsymbol{\lambda}_{z_i,t} - \boldsymbol{\lambda}_{z_j,t}^{\top}\boldsymbol{\lambda}_{z_j,t}$.

[9] As in equation (2.10), one can also express this time-varying expected excess rate of return in terms of a Jensen's inequality term and a logarithmic risk premium term.

Although equation (2.23) demonstrates that tight restrictions related to the prices of risks and the forecasts of excess currency returns implied by the CP factors arise when everything is observable, these return forecasting variables may also enter the determination of the prices of risks, $\boldsymbol{\lambda}_{z_i,t}$ and $\boldsymbol{\lambda}_{z_j,t}$, or they may simply be correlated with the variables that drive these prices of risks, which are not observed, in which case OLS regression of the excess rate of return on $x_{i,t}$ and $x_{j,t}$ does not isolate the pure effect of these variables that arises strictly from the fact that they are the determinants of the prices of the term structure risks, $\boldsymbol{\lambda}_{i,t}$ and $\boldsymbol{\lambda}_{j,t}$. Any restrictions arising from an ATSM specification of $\boldsymbol{\lambda}_{i,t}$ and $\boldsymbol{\lambda}_{j,t}$ are lost in the general regression specification in equation (2.24) because the determinants of $\boldsymbol{\lambda}_{z_i,t}$ and $\boldsymbol{\lambda}_{z_j,t}$ are not included in the regression. Table 2.5 presents the estimated coefficients for equation (2.24) with their asymptotic standard errors in parenthesis for the two sample periods.[10]

Here, we only present the regressions associated with forecasts of the USD excess rates of return from investments in the eight different currencies relegating the results of the remaining 28 non-USD, non-redundant, cross-currency excess rate of return forecasting regressions to the Online Appendix.[11] The statistical significance of the estimates of the coefficients is quantified with three different tests. The $\chi^2(2)$ statistic tests the null hypothesis that ϕ_1 and ϕ_2 equal zero, which tests whether the USD CP factor and the currency j CP factor have forecasting power for the USD excess rate of return on an investment in the currency j money market. The first $\chi^2(1)$ statistic tests the null hypothesis that ϕ_3 equals zero, which tests whether the interest differential has forecasting power; and the second $\chi^2(1)$ statistic tests the null hypothesis that ϕ_3 equals one.

This latter hypothesis is motivated by the typical finding in tests of uncovered interest rate parity that the slope coefficient in a regression of the rate of appreciation of the USD relative to currency j, $\Delta s_{1j,t+1}$, on the interest differential, $r_{j,t} - r_{1,t}$, often produces estimated coefficients that

[10] Appendix C derives the standard errors of the parameters in equation (2.24). These standard errors allow for the fact that the forecasting variables are estimated in first stage regressions.

[11] The findings in these investigations are quite similar to the USD results included in the paper and are consequently not discussed here. Also, note again that for completeness we present the results for the NOK, but because the first sample for the NOK is particularly short, we do not interpret them.

Table 2.5: Forecasts of USD Excess Rates of Return in Currency Markets

CUR j	ϕ_0	ϕ_1	ϕ_2	ϕ_3	$\chi^2(2)$ $\phi_1,\phi_2=0$	$\chi^2(1)$ $\phi_3=0$	$\chi^2(1)$ $\phi_3=1$	R^2
GBP1	-2.92	0.24	-0.24	1.86	0.26	4.31	0.92	0.09
	(3.08)	(0.62)	(1.20)	(0.89)	⟨0.88⟩	⟨0.04⟩	⟨0.34⟩	
GBP2	8.60	-1.89	-2.84	-3.68	2.93	3.57	5.78	0.35
	(5.79)	(1.22)	(1.91)	(1.95)	⟨0.23⟩	⟨0.06⟩	⟨0.02⟩	
EUR1	-1.45	1.09	1.28	1.69	2.75	3.70	0.62	0.16
	(3.56)	(0.73)	(1.35)	(0.88)	⟨0.25⟩	⟨0.05⟩	⟨0.43⟩	
EUR2	8.95	-1.93	-4.21	-0.87	1.98	0.22	1.02	0.18
	(8.03)	(2.26)	(3.59)	(1.85)	⟨0.37⟩	⟨0.64⟩	⟨0.31⟩	
CHF1	3.07	0.09	-0.49	2.47	0.07	7.71	2.73	0.20
	(2.92)	(1.98)	(1.90)	(0.89)	⟨0.97⟩	⟨0.01⟩	⟨0.10⟩	
CHF2	1.71	0.50	-1.22	0.14	1.65	0.01	0.57	-0.00
	(3.20)	(1.24)	(0.97)	(1.15)	⟨0.44⟩	⟨0.90⟩	⟨0.45⟩	
CAD1	0.52	0.13	-0.90	1.46	2.34	3.87	0.38	0.24
	(2.52)	(0.97)	(0.60)	(0.74)	⟨0.31⟩	⟨0.05⟩	⟨0.54⟩	
CAD2	2.22	-3.43	3.17	-3.88	5.63	1.54	2.44	0.35
	(3.96)	(1.59)	(2.44)	(3.12)	⟨0.06⟩	⟨0.21⟩	⟨0.12⟩	
JPY1	10.75	3.13	-2.84	3.94	9.11	17.39	9.68	0.58
	(5.17)	(1.17)	(1.31)	(0.94)	⟨0.01⟩	⟨0.00⟩	⟨0.00⟩	
JPY2	-8.59	0.65	12.01	0.41	8.40	0.11	0.22	0.17
	(4.89)	(2.27)	(6.29)	(1.25)	⟨0.01⟩	⟨0.74⟩	⟨0.64⟩	
AUD1	-3.26	0.24	-0.57	2.50	0.27	9.10	3.27	0.23
	(4.13)	(2.12)	(1.18)	(0.83)	⟨0.87⟩	⟨0.00⟩	⟨0.07⟩	
AUD2	9.62	-2.18	-1.66	-0.68	0.89	0.05	0.32	0.05
	(10.69)	(2.51)	(3.16)	(2.96)	⟨0.64⟩	⟨0.82⟩	⟨0.57⟩	
SEK1	-1.03	-0.76	0.60	0.90	0.16	0.49	0.01	0.03
	(5.64)	(2.61)	(1.72)	(1.29)	⟨0.92⟩	⟨0.49⟩	⟨0.94⟩	
SEK2	7.00	-2.34	-2.61	-1.52	2.19	0.81	2.23	0.16
	(7.00)	(2.80)	(2.09)	(1.69)	⟨0.33⟩	⟨0.37⟩	⟨0.14⟩	
NOK1	-9.92	2.54	0.24	4.91	5.78	11.24	7.12	0.70
	(5.13)	(1.16)	(1.48)	(1.46)	⟨0.06⟩	⟨0.00⟩	⟨0.01⟩	
NOK2	1.47	-2.28	2.55	-1.90	0.65	0.69	1.61	0.11
	(6.21)	(2.83)	(6.88)	(2.29)	⟨0.72⟩	⟨0.41⟩	⟨0.20⟩	

(*Continued*)

Description: The Table presents estimation results for the regression

$$-\Delta s_{1j,t+1} + r_{j,t} - r_{1,t} = \phi_0 + \phi_1 x_{1,t} + \phi_2 x_{j,t} + \phi_3 \left(r_{j,t} - r_{1,t}\right) + \epsilon^s_{1j,t+1},$$

where the dependent variable is the excess USD rate of return on an annual investment in the money market of currency j. The regressors are the CP factors, the fitted return forecasting variables from the term structure regressions for the USD and for currency j, and the difference in the one-year yields between currency j and the USD. Standard errors are in parentheses, and p-values are in angled brackets. The $\chi^2(2)_{\phi_1,\phi_2=0}$ statistic tests the null hypothesis that ϕ_1 and ϕ_2 equal zero. The $\chi^2(1)_{\phi_3=0}$ and $\chi^2(1)_{\phi_3=1}$ statistics test the null hypothesis that ϕ_3 equals either zero or one, respectively. The results for the first sample period are labeled CUR1, and the dependent variables all end in 2003:12. The first samples begin in 1974:12 for the GBP and the EUR; in 1989:03 for the CHF; in 1987:03 for the CAD; in 1986:03 for the JPY; in 1988:04 for the AUD; in 1988:03 for the SEK; and in 1999:03 for the NOK. The results for the second sample period are labeled CUR2, and the dependent variables all begin in 2004:12 and end in 2016:12.

Interpretation: For both samples we are generally unable to reject the null hypothesis that USD and foreign currency CP factors are not significant predictors of the excess USD rate of return on foreign money market investments. Consistent with earlier literature on uncovered interest rate parity, in the first sample we generally strongly reject the null hypothesis that excess returns in foreign money markets are unpredictable using interest differentials. In the post-2003 sample, we are generally unable to reject the null hypothesis that excess returns in foreign money markets are unpredictable using interest differentials. These latter results are consistent with the post-2007 deterioration of returns on foreign currency carry trades.

are significantly negative. That is, the typical Fama (1984) regression is

$$\Delta s_{1j,t+1} = \alpha + \beta \left(r_{j,t} - r_{1,t}\right) + \epsilon^s_{1j,t+1}, \tag{2.25}$$

and the null hypothesis of uncovered interest rate parity is $\beta = 1$. Estimates of β are typically negative. Our specification multiplies this regression by -1, adds the interest differential to both sides, and includes two other variables. Thus, the relation between the two slope coefficients is $-\beta + 1 = \phi_3$. The historical finding that estimated β's in regression (2.25) are negative translates into $\phi_3 > 1$ in our analysis.

For the first sample, most of the coefficients on the CP factors are smaller than their standard errors, and only for the tests associated with the JPY do we find sufficiently large test statistics to reject, at the .05 marginal level of significance or smaller, the null hypothesis that the USD CP factor and the foreign currency CP factor are not significant determinants of the excess rates of return on investments in the foreign money markets.

In contrast, for the first sample, we can reject that the ϕ_3 coefficients are equal to zero at least at the .05 marginal level of significance for all currencies except the SEK. The ϕ_3 estimates are systematically larger than one, but we are only able to reject that they equal one for the JPY at less than the .01 level and for the AUD at the .07 level. The adjusted R^2's for several of the currencies are also substantial. For the currencies with significant estimates of ϕ_3, the R^2's range from .09 for the GBP to .58 for the JPY.

These results are completely consistent with the literature on the FX carry trade, which is a strategy that borrows low interest rate currencies and lends high interest rate currencies. The dependent variable in the regressions is the USD return on the carry trade when $r_{j,t} > r_{1,t}$, and the highly positive values of the slope coefficients indicate that expected USD carry trade profits are conditionally high when $r_{j,t} - r_{1,t}$ is conditionally high.[12]

How does the model do in the post-2003 sample? The answer is not particularly well. The results for the USD excess returns are also presented in Table 2.5. Only for the CAD and the JPY are the $p-$values of the $\chi^2(2)$ statistic testing the significance of the CP factors smaller than .06. The point estimates of ϕ_3 become negative for six of the currencies, and we are generally unable to reject that ϕ_3 equals zero. One exception is the GBP which now has a significantly negative ϕ_3. The fact that many of the coefficients on the CP factors and the interest differentials change signs across the two samples clearly supports the conclusion that the second sample containing the financial crisis is quite different from the pre-2004 sample. These results are also consistent with the post-2007 deterioration in the returns to the carry trade for major currencies.

2.5 Out-of-Sample Results

The previous sections examined the predictability of excess returns in bond and foreign exchange markets with classical asymptotic distribution theory. Inoue and Kilian (2005) argue that such an approach is actually more powerful than out-of-sample experiments, yet such experiments are routinely done and are considered to be a good indicator of instability in the underlying forecasting model. This section consequently examines whether the Cochrane and Piazzesi (2005)

[12]See Daniel, Hodrick, and Lu (2017) for a recent review of the literature on the risks of the carry trade at the monthly holding period horizon. Lustig, Stathopoulos, and Verdelhan (2017) find that investing in the carry trade with longer term bonds while maintaining the one-month holding period is unattractive as the term premiums offset the currency premiums.

model can forecast the excess rates of return in bond and foreign currency markets out of sample. As above, we use the sample period that would have been available when the original paper was written as the in-sample period and treat the post-2003 sample beginning in January 2004 and ending in December 2016 as the out-of-sample period.[13]

We follow Welch and Goyal (2007) and Campbell and Thompson (2007) in assessing the models' out-of-sample forecasts by examining two statistics. The first is the R^2 that compares the mean squared error of the conditional forecasts of excess returns from the term structure model to the mean squared error from assuming that the conditional forecasts of the excess returns are the conditional sample means using data up to that point in time. Analytically, if $\hat{r}_t$ represents the $t-$th out-of-sample forecast from the Cochrane and Piazzesi (2005) model using parameters estimated with all the historical data available at that time, and if $\overline{r}_t$ represents the analogous forecast from the historical sample mean, using the same sample period, then with T_{os} total out-of-sample observations, the mean squared error from the CP forecasts is

$$MSE_{CP} = \frac{1}{T_{os}} \sum_{t=1}^{T_{os}} (r_t - \hat{r}_t)^2 \,, \tag{2.26}$$

and the mean squared error from the historical mean forecasts is

$$MSE_{HM} = \frac{1}{T_{os}} \sum_{t=1}^{T_{os}} (r_t - \overline{r}_t)^2 \,. \tag{2.27}$$

The R^2 is then defined as

$$R^2 = 1 - \frac{MSE_{CP}}{MSE_{HM}} \tag{2.28}$$

The second closely related statistic is the Clark and McCracken (2005) $MSE - F$ which tests for

[13]Because the expected one-year excess return on the two-year bond can be written as $E_t(rx^{(2)}_{t+1}) = -E_t(y^{(1)}_{t+1} - y^{(1)}_t) + (f^{(2)}_t - y^{(1)}_t)$, if the forward spread fails to predict the excess bond return, it must predict the change in the short-term rate. We thank John Cochrane for reminding us of this important caveat which is discussed in Cochrane and Piazzesi (2006).

the equality of the two forecasts:

$$MSE - F = T_{os}\frac{MSE_{HM} - MSE_{CP}}{MSE_{CP}}. \tag{2.29}$$

Panel A of Table 2.6 presents the out-of-sample forecasts of excess bond returns for the Cochrane and Piazzesi (2005) model estimated separately for each currency. The first forecast is 2004:01, and the last is 2016:12.

The results are quite mixed. The model's forecasts are worse than the forecasts based on the historical mean at all maturities for the USD, the EUR, the JPY, the AUD, and the NOK. Only for the GBP do the model forecasts beat historical mean forecast for all maturities. For the CHF, the CAD, and the SEK, the results are mixed across maturities.

To better understand the inability of the forecasts from the estimated Cochrane and Piazzesi (2005) models to beat the forecasts from the historical means, the solid line in Figure 1 presents the evolution of the difference in the cumulative sum of squared errors (SSE) from the forecasts based on the historical mean relative to the CP models' forecasts for the excess returns on the five-year USD, EUR, and JPY bonds.[14]

An increase in the cumulative SSE implies that the CP model forecasts performed better in a given period. The shaded regions show two standard error bands for the SSE and begin after the first 24 months of the out-of-sample period. The dotted line shows in-sample forecasting results from estimation over the full sample.

One sees that for the USD and the EUR, the forecasts of the model were equal to or marginally better than the forecasts from the historical means until 2007 for the USD and 2008 for the EUR, and then both models experienced a deterioration in forecasting ability during the global financial crisis. The model forecasts for the JPY were immediately inferior to those of the historical mean but were recovering in the later part of the sample. Although the point estimates clearly do not favor the models' forecasts, given the volatility of excess returns, we are generally unable to reject the hypothesis that the forecasts are the same. Finally, the dotted lines indicate that in-sample

[14]Figures for all maturities and bonds denominated in other currencies are available in the Online Appendix.

Table 2.6: Out-of-Sample Forecasts for Excess Bond Returns from Cochrane-Piazzesi Models

	R^2				MSE-F			
CUR	$rx_{t+1}^{(2)}$	$rx_{t+1}^{(3)}$	$rx_{t+1}^{(4)}$	$rx_{t+1}^{(5)}$	$rx_{t+1}^{(2)}$	$rx_{t+1}^{(3)}$	$rx_{t+1}^{(4)}$	$rx_{t+1}^{(5)}$
Panel A: Basic Models vs. Historical Means								
USD	-1.51	-1.95	-1.90	-1.68	-87.26	-95.86	-94.97	-90.85
GBP	0.06	0.10	0.11	0.09	9.29	16.50	18.14	15.01
EUR	-0.10	-0.17	-0.25	-0.31	-13.20	-20.65	-28.77	-34.07
CHF	-0.16	0.05	0.05	0.01	-19.75	7.64	7.99	0.73
CAD	0.01	0.00	-0.06	-0.15	1.36	0.62	-7.59	-19.13
JPY	-0.23	-0.13	-0.17	-0.21	-26.65	-16.27	-20.58	-24.94
AUD	-2.27	-2.25	-2.20	-2.08	-100.63	-100.37	-99.73	-97.85
SEK	-0.13	-0.06	0.02	0.07	-16.61	-8.11	2.98	11.26
NOK	-0.97	-1.09	-1.07	-0.97	-71.47	-75.62	-74.86	-71.41
Panel B: Models With Free Constants vs. Historical Means								
USD	-0.66	-1.16	-1.23	-1.15	-57.91	-77.90	-79.97	-77.67
GBP	-0.05	-0.08	-0.11	-0.14	-6.73	-10.30	-14.83	-18.12
EUR	-0.11	-0.15	-0.26	-0.36	-13.80	-19.23	-29.79	-38.44
CHF	-0.31	0.04	0.11	0.08	-34.43	6.61	17.40	13.27
CAD	-0.22	-0.15	-0.16	-0.23	-25.94	-18.55	-20.13	-27.27
JPY	-0.21	-0.10	-0.12	-0.11	-25.21	-13.39	-16.08	-14.86
AUD	-1.91	-2.15	-2.44	-2.56	-95.18	-99.03	-102.88	-104.30
SEK	-0.11	-0.12	-0.10	-0.08	-14.05	-15.73	-13.21	-11.14
NOK	-0.61	-0.83	-0.98	-1.06	-55.20	-65.79	-71.84	-74.56

Description: The Table reports two statistics that compare the out-of-sample forecasts from recursive estimations of two versions of the Cochrane and Piazzesi (2005) model for the excess rates of returns on bonds denominated in different currencies compared to the forecasts based only on the historical mean excess rates of return. Panel A contains results for the basic model, and Panel B contains results for models estimated with free constant terms. The first statistic is the R^2, which is calculated as one minus the ratio of the mean squared error of the CP forecasts to the mean squared error of the historical mean. The second statistic tests the equality of the forecasts and is the Clark and McCracken (2005) $MSE - F$ statistic. The sample periods for the dependent variables during the initial in-sample estimation all end in 2003:12, which is the end of the Cochrane and Piazzesi (2005) sample. The samples begin in 1974:12 for the USD, the GBP, and the EUR; in 1989:03 for the CHF; in 1987:03 for the CAD; in 1986:03 for the JPY; in 1988:04 for the AUD; in 1988:03 for the SEK; and in 1999:03 for the NOK. The out-of-sample periods are all 2004:01 to 2016:12.

Interpretation: For all currencies and for both types of the Cochrane and Piazzesi (2005) models, the out-of-sample forecasts of excess bond returns from the models are inferior to the forecasts from the historical means of excess returns.

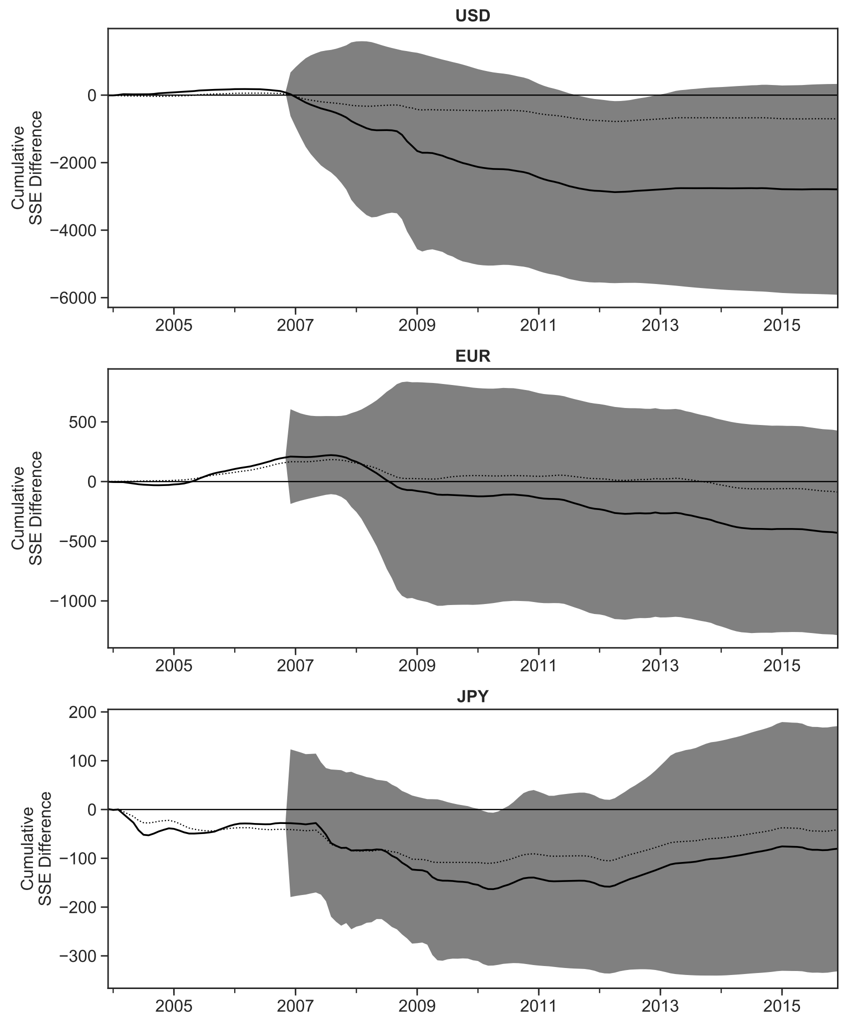

Figure 2.1: Out-of-Sample Results

Description: The figures plot the differences in the out-of-sample cumulative sums of squared errors from forecasts based on the historical means and the Cochrane-Piazzesi models for excess returns on the five-year USD, EUR, and JPY bonds. An increase in the cumulative SSE (the solid line) implies that CP model performed better in a given period. The shaded regions show two standard error bands for the SSE. The dotted line shows in-sample results from estimation over the full sample.

Interpretation: The Figures demonstrate that the relative deterioration in the model forecasts for the USD and the EUR begins in 2007 and 2008 with the onset of the global financial crisis. While the point estimates favor the historical means, the large standard errors remind us that distinguishing between out-of-sample forecasts is inherently difficult.

forecasts of the models from a constant set of parameters were generally inferior to the historical mean during this period.

2.5.1 An Alternative Model with Free Constants

In discussing the relation of the Cochrane and Piazzesi (2008) ATSM to the empirical model in Cochrane and Piazzesi (2005), we noted that the former does not constrain the constant terms to have the same factor of proportionality across maturities as is imposed by the latter. To see whether relaxing this constraint which formally nests the historical mean model as a constrained version of the larger model, we recursively estimated the model with free constant terms for each maturity.[15]

The results of these out-of-sample forecasts are presented in Panel B of Table 2.6. All of the R^2's except for maturities 3, 4, and 5 for the CHF are negative. We conclude that the forward spreads, when used in this way, provide no useful out-of-sample forecasting power for the excess bond returns.

2.5.2 Constraining Parameters Across Currencies

In out-of-sample forecasting situations, it is often advised to limit the number of free parameters that are estimated. We experimented with this Occam's razor intuition and recursively estimated the free constant model, described in the previous subsection, subject to the restriction that the γ parameters are the same for all countries. The out-of-sample forecasting results for this constrained model are better than the historical mean for the CAD, the GBP, and the JPY, but worse for the other currencies and for all maturities. These results are consequently presented in the Online Appendix.

[15] The equations of the free constant model are described in the Online Appendix which also reports the parameter estimates. These estimates do not different substantively from the estimates reported in the paper for the basic model.

2.5.3 Evolution of the USD Parameters

The failure of the model in the out-of-sample forecasting experiments and the rejection of equality of coefficients across sub-periods suggests substantial parameter instability. While a full analysis of this issue is not something we have space to accomplish in this article, Figure 2.2 presents the recursive estimates of the parameters of the Cochrane and Piazzesi (2005) model for the USD term structure as they evolve in the out-of-sample estimation period.[16] The estimates of the b_n parameters remain incredibly stable as the four lines are virtually horizontal.

It is also clear that beginning in 2008 with the advent of the financial crisis, the estimated $\gamma(2)$ changes over the course of two years from positive to negative, the estimated $\gamma(3)$ begins a slow decline, and the estimated $\gamma(5)$ experiences a steady increase. The estimated $\gamma(4)$ is reasonably constant after a blip in 2009. Because these are recursive estimates that use all of the sample to that point in time, they are more stable than would be recursive rolling estimates that use the same sample size at each point in time. In that sense, the slow evolution masks more dramatic changes.

2.5.4 Out-of-Sample Forecasts of Currency Returns

Give the inability of the CP factors to forecast excess rates of return in the currency markets reported in Table 2.5 and the changes in the signs of the estimated coefficients on the interest differentials in the in-sample regressions, the reader should expect that this specification will not be useful in out-of-sample of the excess return. For completeness, we present these results in Table 2.7.

The results are indeed as anticipated as the out-of-sample forecasts from the model are unable to beat the historical mean excess returns of all currencies versus the USD.

[16] Comparable Figures for the other currencies for both the basic model and the free constant model are available in the Online Appendix.

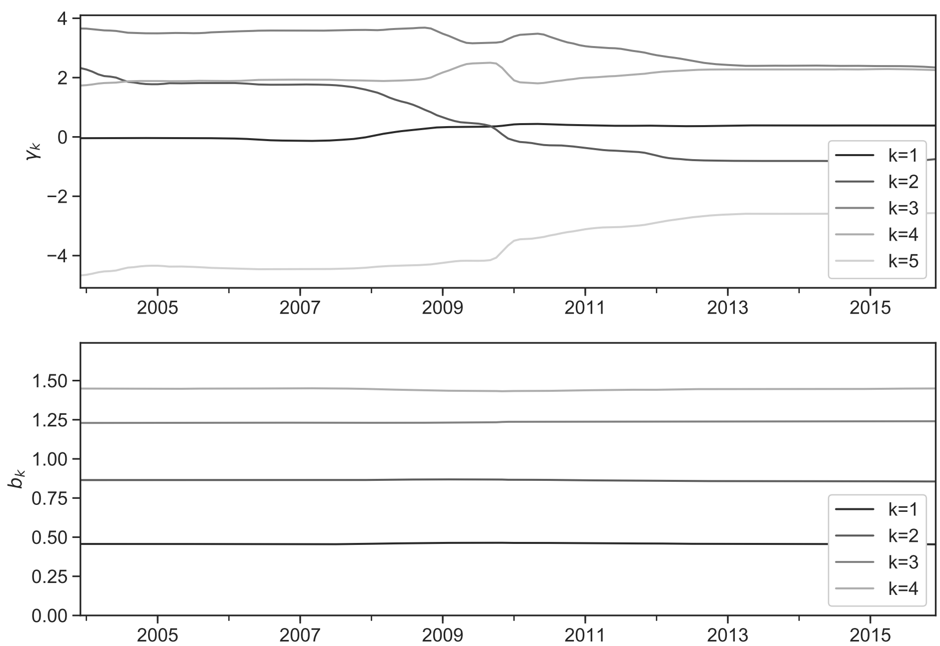

Figure 2.2: Evolution of the USD Parameters

Description: The figures present the evolution of the estimated $\gamma(k)$ parameters in the top and the $b(k)$ parameters in the bottom as the sample is extended from 2003:12 to 2016:12.

Interpretation: The stability of the $b(k)$ parameter estimates is remarkable while it is clear that the instability of the $\gamma(k)$ parameters is mostly related to the period between 2008 and 2013 after which the coefficients appear to have stabilized.

Table 2.7: Out-of-Sample Forecasts for Excess Foreign Exchange Returns: Cochrane-Piazzesi Factors and Interest Differentials vs. Historical Means

CUR	R^2	MSE-F
GBP	-0.08	-10.61
EUR	-0.25	-28.72
CHF	-0.18	-21.92
CAD	-0.16	-19.96
JPY	-0.43	-43.34
AUD	-0.14	-17.26
SEK	-0.02	-3.31
NOK	-0.33	-35.61

Description: The Table reports two statistics that compare the out-of-sample forecasts from recursive estimation of equation (2.24) for the excess return in USD on one-year investments in the money markets of different currencies to the forecasts based on the historical mean excess rates of return on those currencies. The first statistic is the R^2, which is calculated as one minus the ratio of the mean squared error of the CP forecasts to the mean squared error of the historical mean. The second statistic tests the equality of the forecasts and is the Clark and McCracken (2005) $MSE - F$ statistic. The sample periods for the dependent variables during the initial in-sample estimation all end in 2003:12, which is the end of the Cochrane and Piazzesi (2005) sample. The samples begin in 1974:12 for the USD, the GBP, and the EUR; in 1989:03 for the CHF; in 1987:03 for the CAD; in 1986:03 for the JPY; in 1988:04 for the AUD; in 1988:03 for the SEK; and in 1999:03 for the NOK. The out-of-sample periods are all 2004:01 to 2016:12.

Interpretation: Out-of-sample forecasts of currency returns using the Cochrane-Piazzesi forecasting factors and interest differentials are generally inferior to forecasts from the historical means of the currency returns.

2.6 Related Literature

The Cochrane and Piazzesi (2005, 2008) papers generated a vast literature. In this section we briefly review what we consider to be the most important contributions in the literature that are related to our paper.[17]

Dahlquist and Hasseltoft (2013, 2016) and Sekkel (2011) were the first to extend the Cochrane and Piazzesi (2005) model to the bond markets of additional currencies. Dahlquist and Hasseltoft (2013) examine the bond markets of the USD, the CHF, the EUR, and the GBP, as well as examining the USD returns on the foreign bonds. They use a sample period from January 1975 to December 2009, and the CP factor is constructed from projections onto the five forward rates as in the original paper. They estimate local currency CP factors, and they also construct a global CP factor as a GDP-weighted average of the local CP factors. They find that the global CP factor provides some additional explanatory power relative to the local CP factors. In Dahlquist and Hasseltoft (2016) they extend their analysis adding the bond markets of the AUD, the CAD, the DKK, the JPY, the NOK, and the SEK; and they employ a sample period from December 1999 to December 2013. They find support for the model in all currencies, but they do not investigate the stability of the coefficients. We similarly find that the model works well in our two non-overlapping samples, but we are able often able to reject that the parameters are the same in the two samples.

Wright (2011) examines the term structures of interest rates for the G-10 countries by estimating ATSMs as in Joslin, Priebsch, and Singleton (2014). He studies the implied risk premiums or term premiums, defined as the difference between the long-term yields and expectations of future spot interest rates, finding that these term premiums have generally declined in most countries over the sample period from January 1990 to May 2009. Bauer, Rudebusch, and Wu (2014) dispute these conclusions noting that after correcting for small sample bias in the coefficient estimates, the term premiums show a pronounced countercyclical pattern as was found by Cochrane and Piazzesi

[17] Because the Cochrane and Piazzesi (2005, 2008) papers have 1,400 and 328 Google Scholar citations, respectively, as of March 2019, our literature review must be highly selective.

(2005). We have not attempted to bias-correct our coefficient estimates, and doing so could affect our conclusions.

Sekkel (2011) uses the Wright (2011) data to estimate the Cochrane and Piazzesi (2005) model, but he projects the excess returns only onto the one, three, and five year forward rates. He finds that the performance of the model deteriorates during the global financial crisis. We use averages of four forward spreads, and we did not experiment with leaving out the two-year spread. We doubt this affects our results, but the financial crisis clearly strongly influences our findings.

Consistent with the finding of Cochrane and Piazzesi (2005) that the CP factor is not spanned by the first three principal components of bond yields, Duffee (2011) documents that almost half of the variation in U.S. dollar (USD) bond risk premiums cannot be detected using the cross-section of yields. He finds that fluctuations in this hidden component have strong forecasting power for both future short-term interest rates and excess bond returns. The hidden component is negatively correlated with aggregate economic activity, but macroeconomic variables explain only a small fraction of variation in the hidden factor.

Koijen, Lustig, and Van Nieuwerburgh (2017) model the stochastic discount factor as depending on the Cochrane and Piazzesi (2005) forecasting factor as well as the return on the stock market and the level of the term structure of interest rates. They demonstrate that such a model does well in simultaneously pricing returns on value and growth stocks in additional to USD zero-coupon bonds. We suspect that the results in Koijen, Lustig, and Van Nieuwerburgh (2017) are affected by the instability in the CP factors that we document, but we have not investigated how their model would perform in the different samples we investigate.

Kessler and Scherer (2009), Thornton and Valente (2012), Zhu (2015), and Sarno, Schneider, and Wagner (2016) perform out-of-sample forecasting analyses with the Cochrane and Piazzesi (2005) model. Kessler and Scherer (2009) assess the performance of trading strategies based on a one-month forecast horizon using data from seven currencies (the AUD, CAD, CHF, EUR, GBP, JPY, and the USD) for the sample period February 1997 to July 2007. They use either a 36 or 60 month rolling window to estimate the parameters of the forecasting equation implying that they have either 88 or 64 true out-of-sample forecasts. They find slightly positive but only marginally significant trading profits. Our out-of-sample results are for an annual holding period and show no evidence of useful predictability.

Thornton and Valente (2012) investigate the out-of-sample predictability of USD bond excess returns and assess the economic value of the forecasting ability of empirical models based on Fama and Bliss (1987) and Cochrane and Piazzesi (2005). Their results show that the information content of forward rates does not generate systematic economic value to investors in a dynamic asset allocation exercise. Furthermore, they find that the models do not outperform the no-predictability benchmark, and their relative performance deteriorates over time. We find similar results for the bond markets of other currencies and for the foreign currency markets.

Zhu (2015) explores the forecasting ability of a global CP factor constructed as the forecast of the average returns on the two through five year maturity bonds averaged over four currencies (the EUR, JPY, GBP, and USD) when regressed on the four individual currency CP factors. The full sample period is January 1980 to December 2011, and the out-of-sample period begins in January 1992. In contrast to our findings, Zhu (2015) finds statistically significant out-of-sample forecasts that beat the historical mean return for all four countries. This is true even when only the local currency CP factors are included in the analysis and when the out-of-sample period is restricted to the global financial crisis, 2008-2011. Because these results are so inconsistent with ours, we tried to replicate some of the findings in Zhu (2015). Although we were able to match his summary statistics with our data, we were not able to replicate his out-of-sample findings.

Sarno, Schneider, and Wagner (2016) investigate ATSMs for the USD bond market and find that their implied time-varying risk premiums do not provide important increases in utility to investors over and above inferences about expected future spot interest rates implied by the expectations hypothesis of the term structure with constant risk premiums. While we find strong in-sample evidence of time-varying risk premiums in our two samples, the changing nature of the parameters and the lack of out-of-sample predictability is consistent with their findings.

Turning to the international implications of the modeling, Sarno, Schneider, and Wagner (2012) find that separately estimated ATSMs for two currencies, both of which provide very small pricing errors for zero-coupon bonds denominated in those currencies, are not highly correlated with the relative rate of appreciation of those currencies in the foreign exchange market. These results clearly support our approach to testing the international implications of the Cochrane and Piazzesi (2008) model that allows for risks in the general stochastic discount factor that are not priced in the bond markets.

Our results are also related to the vast literature examining the uncovered interest rate parity (UIRP) hypothesis. Although Chinn and Meredith (2004) provide support for UIRP at the annual horizon, our results are more consistent with the conclusions of Bekaert, Wei, and Xing (2007), who argue that UIRP is violated at longer horizons just as is typically the case at the shorter monthly horizon, although here too we find evidence of parameter instability.

2.7 Conclusions

In this paper we document substantive instabilities in the empirical analysis of risk premiums in bond and foreign exchange markets. There are many directions that research on time varying risk premiums in bond and foreign exchange markets could go. Here we review some recent approaches.

One puzzle we uncover is the observation that there is a strong one-factor structure to the forecasts of expected returns in the bond markets in a particular sample of data, but a different one-factor structure in another sample. Modeling the sources of these structural changes should be high on the research agenda of fixed income research. It is also puzzling that the CP factors in two currencies have such strong predictability in their respective bond markets but not in the foreign exchange market between the two currencies. Investigating why these markets are not more closely linked should also be on the research agenda.

In extending the Cochrane and Piazzesi (2005) model to additional currencies and considering its international implications, we have not addressed the term structure literature arguing that macroeconomic variables, such as inflation and employment, have additional forecasting power over and above that available in bond yields. In this regard we cite two recent critiques of this literature. First, Ghysels, Horan, and Moench (2018) find that several studies touting the significantly improved forecasting performance of macroeconomic variables above that provided by yields overstate their importance because the studies use revised data. Ghysels, Horan, and Moench (2018) find that use of real time U.S. data substantially reduces the implied predictive power. Second, Bauer and Hamilton (2017) argue that after taking account of small sample distortions in the test statistics induced by the use of macro variables with trends, the evidence for additional predictability from macro variables is much weaker.

It is interesting though that Jotikasthira, Le, and Lundblad (2015) document highly correlated yield curve fluctuations across different currencies. They argue that common macroeconomic shocks influence bond yields both through a monetary policy channel and through a risk compensation channel. Using data from the U.S., the UK, and Germany, they find that world inflation and the level of the U.S. yield curve explain over two-thirds of the covariation of yields at all maturities and that these effects operate largely through the risk compensation channel for long-term bonds.

In a related finding, Pericoli and Taboga (2012) propose a two-country no-arbitrage term-structure model to analyze the joint dynamics of bond yields, macroeconomic variables, and the exchange rate. The model demonstrates how exogenous shocks to the exchange rate affect the yield curves, how bond yields co-move in different countries and how the exchange rate is influenced by interest rates, macroeconomic variables and time-varying bond risk premiums. Upon estimating the model with U.S. and German data, they find that time-varying bond risk premiums account for a significant portion of the variability of the exchange rate. As we mentioned above, the correlations we find between the CP factors denominated in different currencies are certainly suggestive that global risk premiums are driven by common international investors, but we leave this issue for future research because addressing these issues in our multiple currency context is beyond what can be accomplished in a given article.

As in the original analysis of Cochrane and Piazzesi (2005), we have focused exclusively on the annual forecasting horizon. Most bond market ATSMs are estimated at the monthly horizon and typically find that monthly risk premiums are driven by more than one state variable. In contrast, we find the strong one-factor structure originally documented by Cochrane and Piazzesi (2005) at the annual horizon. Examining the dynamics of the state variables in monthly models and seeing whether they imply a single state variable at the annual horizon would be an interesting project. Recent papers that examine multiple horizons include Bacchetta and Van Wincoop (2010), Engel (2016), Lustig, Stathopoulos, and Verdelhan (2017), and Chernov and Creal (2018) who find interesting patterns in expected returns at different horizons.

We also have only employed data on bonds with a maximum of five years to maturity as in the original paper of Cochrane and Piazzesi (2005). We treat these data sets as providing yields on actual bonds, as in most of the term structure literature. In contrast, Pancost (2018) uses the raw price data on all outstanding U.S. Treasury bonds rather than fitted zero-coupon yield curves constructed either by the bootstrap method of Fama and Bliss (1987) or the functional form approaches of Nelson and Siegel (1987) or Svensson (1995). Pancost (2018) finds that the two methods of fitting yield curves do well for maturities less than five years, but during the financial crisis there appears to be two separate yield curves in actual data for bonds with maturities between five years and 12 years corresponding to whether the bonds have been outstanding for more or less than 15 years. He also finds that prices of risk estimated from vector autoregressions of bond market factors do not forecast the returns on actual Treasury bonds.

While ATSMs are typically developed under the assumption of rational expectations, it may be the case that behavioral finance with its time varying sentiments could be responsible for our findings of model instability. One recent empirical analysis with a behavioral slant is Brooks and Moskowitz (2017), who examine quarterly returns on bonds denominated in AUD, CAD, EUR, GBP, JPY, SEK, and USD. Using panel data methods with time fixed effects, they argue that measures of value, carry, and momentum dominate the CP factor in forecasting excess returns. While we are forecasting actual excess returns on a currency by currency basis, the panel data approach of Brooks and Moskowitz (2017) implies that they are forecasting deviations from cross-sectional average returns.

Another alternative approach to these issues would rely on the analysis of Krishnamurthy and Vissing-Jorgensen (2011) who argue that the supply of U.S. Treasury securities affects the level and slope of the yield curve. Do such changes in quantities also affect the risk premiums in the other bond markets and in the foreign exchange markets? Valchev (2017) answers this question affirmatively. It is natural to think that major changes in monetary and fiscal policies, including the quantitative easing done by major central banks during the international financial crisis, could induce the changes in the parameters of the CP factors and the resulting changes in forecasting power that we observe. Actually demonstrating this empirically is a challenging task. Evidence of substantive structural change in the international financial markets can also be found in the deviations from covered interest rate parity documented by Du, Tepper, and Verdelhan (2018) and Rime, Schrimpf, and Syrstad (2017).[18]

Rather than focusing on time varying risk premiums, Valchev (2017) and Jiang, Krishnamurthy, and Lustig (2018) are two recent papers that empirically explore differential time-varying liquidity premiums, or non-pecuniary returns, on government bonds as explanations of rates of currency depreciation.

Engel (2016) notes that countries with high real interest rates tend to have currencies that are stronger than can be accounted for by the path of expected real interest differentials under UIRP. He observes that these two findings have contradictory implications for the relationship of the foreign-exchange risk premium and interest-rate differentials and shows that existing models appear unable to account for both puzzles. He then introduces a model, in which short-term assets can have liquidity premiums as in Nagel (2016a), that potentially reconciles the two sets of findings. Our findings of significant differences in parameters across samples that exclude and include the global financial crisis and the illiquidity observed in many asset markets during the financial crisis as well as the role that government bonds as safe assets play in the financial system are certainly suggestive that more research along these lines is warranted. Ultimately, we need to ask if these liquidity based models are able to explain the instabilities that we document.

[18] Andersen, Duffie, and Song (2019) provide a theoretical explanation for deviations from covered interest rate parity in a world with highly levered, risky financial market makers.

Our econometric analysis also is conducted under the standard assumption that investors have rational expectations and that the data are stationary and ergodic. It has long been recognized that changes in monetary policy regimes can cause problems with econometric analysis of the term structure. Fuhrer (1996) argues that investors are aware of changes in regimes but do not anticipate future changes, which he views as a compromise between full rationality and learning. Bekaert, Hodrick, and Marshall (2001) argue that so-called peso problems, caused by differences between the frequency of realizations of the data and the conditional distributions investors had at the time that they set bond prices, could be responsible for the anomalies observed in the term structure literature.

The necessity for investors to learn about changes in monetary policy, the rate of inflation, or the real interest rate are also important areas of recent research that relaxes the rational expectations assumption. Piazzesi, Salomao, and Schneider (2015) note that professional forecasts of interest rates differ from those based on regressions. They build on the insights of Froot (1989) who argued that evidence against the expectations hypothesis of the term structure was plausibly due to the failure of the rational expectations assumption imposed in the tests rather than to failures of the expectations hypothesis itself.

Giacoletti, Laursen, and Singleton (2016) argue that marginal investors in the bond market act as Bayesian learners to form prospective real-time views about bond market risks. While the sources of risks are the first three principal components of the yield curve, knowledge of the extent of disagreement among professionals is informative about how today's yield curve will impact its future shape and thus the prices of risks.

The studies cited here provide some interesting directions in which research can go. Most of these papers do not investigate time variation in the parameters of their empirical models. Our paper provides a set of challenging empirical results demonstrating more attention should be devoted to this type of analysis. The paper also provides interesting empirical evidence showing an absence of links that should theoretically be present between the term structures of interest rates in two currencies and the currency market between them.

2.8 Appendix A: Data

Data on the term structures of interest rates for the different currencies were obtained from several sources. The USD data are from the CRSP Fama-Bliss database. This is the same source as Cochrane and Piazzesi (2005). For yields from the non-USD term structures, we obtained data from Jonathan Wright's web site. These data were used in Wright (2011). We updated the data from the web sites of the respective central banks. The monthly term structure data all end in December 2016. The data begin in June 1952 for the USD, in January 1970 for the GBP, in January 1973 for the EUR spliced with the Deutsche mark prior to 1999; in March 1988 for the CHF; in January 1986 for the CAD; in January 1985 for the JPY; in February 1987 for the AUD; in January 1987 for the SEK; and in January 1998 for the NOK. The exchange rate data are from the St. Louis Federal Reserve Bank FRED database. The sample period is January 1973 to December 2016 for all exchange rates. The Online Appendix provides additional detail about the exact sources of the data including URLs where the data may be updated. All data and programs are also available in this paper's online *Critical Finance Review* depository.

2.9 Appendix B: The Affine Model Solutions

The solutions to the coefficients of the natural logarithms of the bond prices in the affine model given in equation (2.9) are the following difference equations:

$$
\begin{aligned}
A_n &= A_{n-1} - \delta_0 + \boldsymbol{B}_{n-1}^{\top} (\boldsymbol{\mu} - \boldsymbol{\Sigma} \boldsymbol{\lambda}_0) + (1/2) \boldsymbol{B}_{n-1}^{\top} \boldsymbol{\Sigma} \boldsymbol{\Sigma}^{\top} \boldsymbol{B}_{n-1} && (2.30) \\
\boldsymbol{B}_n^{\top} &= -\boldsymbol{\delta}_1^{\top} + \boldsymbol{B}_{n-1}^{\top} (\boldsymbol{\Phi} - \boldsymbol{\Sigma} \boldsymbol{\lambda}_1) && (2.31)
\end{aligned}
$$

with initial conditions $A_0 = 0$ and $\boldsymbol{B}_0 = \mathbf{0}$. By defining $\boldsymbol{\Phi}^* = (\boldsymbol{\Phi} - \boldsymbol{\Sigma} \boldsymbol{\lambda}_1)$, we can write equation (2.31) as

$$
\boldsymbol{B}_n^{\top} = -\boldsymbol{\delta}_1^{\top} (\mathbf{I} - \boldsymbol{\Phi}^*)^{-1} (\mathbf{I} - \boldsymbol{\Phi}^{*n}) \quad (2.32)
$$

2.10 Appendix C: The Standard Errors

This appendix derives the standard errors for the two-step estimation of the term structure models that generate the CP forecasting factors and the corresponding forecasting equation for the excess return on investing USD in the currency j money market. Let $\overline{\varepsilon}_{1,t+1}$ and $\overline{\varepsilon}_{j,t+1}$ be the error terms in equation (2.5) for the term structure regressions associated with the USD and currency j, respectively. The error term, $\varepsilon^{s}_{j,t+1}$, from the currency market is defined in equation (2.24). Let $\mathbf{h}_{1j,t} \equiv \left[1,\ x_{1,t},\ x_{j,t},\ r_{j,t} - r_{1,t}\right]^{\mathsf{T}}$ represent the vector of regressors in equation (2.24), where $x_{j,t} = \widehat{\boldsymbol{\gamma}}_{j}^{\mathsf{T}}\overline{\mathbf{fs}}_{j,t}$ is the return forecasting variable from the estimation of equation (2.5) for currency j, and let $\boldsymbol{\phi}$ represent the vector of parameters in equation (2.24). Then, the orthogonality conditions associated with the forecasts of the average excess returns in the two bond markets and the excess rate of return in the currency market are

$$\begin{aligned} E\left[\overline{\mathbf{fs}}_{1,t} \cdot \overline{\varepsilon}_{1,t+1}\right] &= \mathbf{0} \\ E\left[\overline{\mathbf{fs}}_{j,t} \cdot \overline{\varepsilon}_{j,t+1}\right] &= \mathbf{0} \\ E\left[\mathbf{h}_{1j,t} \cdot \varepsilon^{s}_{j,t+1}\right] &= \mathbf{0}. \end{aligned} \tag{2.33}$$

The parameter vector is $\boldsymbol{\theta} = \left[\boldsymbol{\gamma}_{1}^{\mathsf{T}},\ \boldsymbol{\gamma}_{j}^{\mathsf{T}},\ \boldsymbol{\phi}^{\mathsf{T}}\right]^{\mathsf{T}}$. Let $\mathbf{g}_{T}(\boldsymbol{\theta})$ denote the sample mean of the orthogonality conditions in the system of equations given in (2.33). Because the system is just identified, these sample orthogonality conditions can be set to zero, and the asymptotic variance of the parameter estimates can be estimated as

$$\mathbf{V}(\boldsymbol{\theta}) = \frac{1}{T}\mathbf{D}_{T}^{-1}\mathbf{S}_{T}\mathbf{D}_{T}^{-1\mathsf{T}} \tag{2.34}$$

where

$$\mathbf{D}_{T} = \frac{\partial \mathbf{g}_{T}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}^{\mathsf{T}}} \tag{2.35}$$

is the sample estimate of the Jacobian of the orthogonality conditions, $\mathbf{D}$, which is defined below, and

$$\mathbf{S_T} \equiv \mathbf{C}_0 + \sum_{k=1}^{K} \frac{K-k}{K} \left(\mathbf{C}_k + \mathbf{C}_k^\intercal \right), \tag{2.36}$$

is the sample estimate of the variance of the orthogonality conditions. The autocovariances are estimated with

$$\mathbf{C}_k \equiv \frac{1}{T} \sum_{t=k+1}^{T} \mathbf{g}_t \mathbf{g}_{t-k}^\intercal \tag{2.37}$$

where $\mathbf{g}_t$ is the vector of observations on the orthogonality conditions are time t, and we use $K = 18$.

The derivatives in equation (2.35) are sample estimates of

$$\mathbf{D} = \begin{bmatrix} -E\left(\overline{\mathbf{fs}}_{1,t}\overline{\mathbf{fs}}_{1,t}^\intercal\right) & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & -E\left(\overline{\mathbf{fs}}_{j,t}\overline{\mathbf{fs}}_{j,t}^\intercal\right) & \mathbf{0} \\ \mathbf{D}_1 & \mathbf{D}_2 & \mathbf{D}_3 \end{bmatrix}$$

where $\mathbf{D}_1 \equiv \nabla_{\gamma_1^\intercal} E\left(\varepsilon_{j,t+1}^s \mathbf{h}_{1j,t}\right)$, $\mathbf{D}_2 \equiv \nabla_{\gamma_j^\intercal} E\left(\varepsilon_{j,t+1}^s \mathbf{h}_{1j,t}\right)$, and $\mathbf{D}_3 \equiv \nabla_{\phi^\intercal} E\left(\varepsilon_{j,t+1}^s \mathbf{h}_{1j,t}\right)$, respectively. We estimate $\mathbf{D}_T$ numerically using Python's numdifftools package.

From the structure of the $\mathbf{D}$ matrix and the partitioned inverse formula, one sees that the variances of the estimates of $\boldsymbol{\gamma}_1$ and $\boldsymbol{\gamma}_j$ are not affected by the estimation of $\boldsymbol{\phi}$ whereas the variances of the latter parameters are affected by the estimation of the former.

Bibliography

Acharya, V. V., T. Philippon, and M. Richardson. 2016. Measuring Systemic Risk for Insurance Companies. In F. Hufeld, R. S. J. Koijen, and C. Thimann, eds., *The Economics, Regulation, and Systemic Risk of Insurance Markets*. Oxford: Oxford University Press.

Adrian, T., E. Etula, and T. Muir. 2014. Financial Intermediaries and the Cross-Section of Asset Returns. *The Journal of Finance* 69:2557–96.

Andersen, L., D. Duffie, and Y. Song. 2019. Funding value adjustments. *The Journal of Finance* 74:145–92.

Aon Benfield. 2009. Insurance-Linked Securities - Adapting to an Evolving Market.

———. 2018. Insurance-Linked Securities - Alternative Capital Fortifies Its Position.

Bacchetta, P., and E. Van Wincoop. 2010. Infrequent portfolio decisions: A solution to the forward discount puzzle. *American Economic Review* 100:870–904.

Backus, D. K., S. Foresi, and C. I. Telmer. 2001. Affine term structure models and the forward premium anomaly. *The Journal of Finance* 56:279–304.

Bansal, R., D. Kiku, and M. Ochoa. 2017. Price of Long-Run Temperature Shifts in Capital Markets. Duke University Working Paper.

Bao, J., J. Pan, and J. Wang. 2011. The Illiquidity of Corporate Bonds. *The Journal of Finance* 66:911–46.

Barberis, N., and M. Huang. 2008. Stocks as Lotteries : The Implications of Probability Weighting for Security Prices. *American Economic Review* 98:2066–100.

Barro, R. J. 2006. Rare Disasters and Asset Markets in the Twentieth Century. *The Quarterly Journal of Economics* 121:823–66.

Bauer, M. D., and J. D. Hamilton. 2017. Robust bond risk premia. *The Review of Financial Studies* 31:399–448.

Bauer, M. D., and G. D. Rudebusch. 2017. Interest rates under falling stars. San Francisco Federal Reserve Bank Working Paper.

Bauer, M. D., G. D. Rudebusch, and J. C. Wu. 2014. Term premia and inflation uncertainty: empirical evidence from an international panel dataset: comment. *American Economic Review*

104:323–37.

Bekaert, G., R. J. Hodrick, and D. A. Marshall. 2001. Peso problem explanations for term structure anomalies. *Journal of Monetary Economics* 48:241–70.

Bekaert, G., M. Wei, and Y. Xing. 2007. Uncovered interest rate parity and the term structure. *Journal of International Money and Finance* 26:1038–69.

Bolton, P., and M. Kacperczyk. 2019. Do Investors Care about Carbon Risk? Columbia Business School Working Paper.

Brandt, M. W., J. H. Cochrane, and P. Santa-Clara. 2006. International risk sharing is better than you think, or exchange rates are too smooth. *Journal of Monetary Economics* 53:671–98.

Brandt, M. W., and P. Santa-Clara. 2002. Simulated likelihood estimation of diffusions with an application to exchange rate dynamics in incomplete markets. *Journal of Financial Economics* 63:161–210.

Braun, A. 2016. Pricing in the Primary Market for Cat Bonds: New Empirical Evidence. *Journal of Risk and Insurance* 83:811–47.

Brooks, J., and T. J. Moskowitz. 2017. Yield curve premia. Yale School of Management Working Paper.

Brunnermeier, M. K., and Y. Sannikov. 2014. A Macroeconomic Model with a Financial Sector. *American Economic Review* 104:379–421.

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* 24:853–91.

Campbell, J. Y., and S. B. Thompson. 2007. Predicting excess stock returns out of sample: Can anything beat the historical average? *The Review of Financial Studies* 21:1509–31.

Chen, H., S. Joslin, and S. X. Ni. 2019. Demand for Crash Insurance , Intermediary Constraints , and Risk Premia in Financial Markets. *The Review of Financial Studies* 32:228–65.

Chernov, M., and D. D. Creal. 2018. Multihorizon currency returns and purchasing power parity. NBER Working Paper No. 25963.

Chinn, M. D., and G. Meredith. 2004. Monetary policy and long-horizon uncovered interest parity. *IMF staff papers* 51:409–30.

Cieslak, A., and P. Povala. 2015. Expected returns in treasury bonds. *The Review of Financial Studies* 28:2859–901.

Clark, T. E., and M. W. McCracken. 2005. Evaluating direct multistep forecasts. *Econometric Reviews* 24:369–404.

Cochrane, J. H. 2015. Comments on "Robust Bond Risk Premia" by Michael Bauer and Jim Hamilton. Manuscript, Stanford University.

Cochrane, J. H., and M. Piazzesi. 2005. Bond risk premia. *American Economic Review* 95:138–60.

———. 2006. Appendix to "bond risk premia". University of Chicago Working Paper.

———. 2008. Decomposing the yield curve. University of Chicago Working Paper.

Cummins, J. D., D. Lalonde, and R. D. Phillips. 2004. The basis risk of catastrophic-loss index securities. *Journal of Financial Economics* 71:77–111.

Dahlquist, M., and H. Hasseltoft. 2013. International bond risk premia. *Journal of International Economics* 90:17–32.

———. 2016. International bond risk premia. In P. Veronesi, ed., *Handbook of Fixed Income Securities*, chap. 9. New York: Wiley.

Daniel, K., R. J. Hodrick, and Z. Lu. 2017. The carry trade: Risks and drawdowns. *Critical Finance Review* 6:211–62.

Daniel, K. D., R. B. Litterman, and G. Wagner. 2018. Applying Asset Pricing Theory to Calibrate the Price of Climate Risk. Columbia Business School Working Paper.

Dick-Nielsen, J. 2009. Liquidity Biases in TRACE. *The Journal of Fixed Income* 19:43–55.

———. 2014. How to clean Enhanced TRACE data. Copenhagen Business School Working Paper.

Dick-Nielsen, J., P. Feldhütter, and D. Lando. 2012. Corporate bond liquidity before and after the onset of the subprime crisis. *Journal of Financial Economics* 103:471–92.

Du, W., A. Tepper, and A. Verdelhan. 2018. Deviations from covered interest rate parity. *The Journal of Finance* 73:915–57.

Duffee, G. R. 2011. Information in (and not in) the term structure. *The Review of Financial Studies* 24:2895–934.

Duffie, D. 2010. Presidential Address: Asset Price Dynamics with Slow-Moving Capital. *The Journal of Finance* 65:1237–67.

Elenev, V., T. Landvoigt, and S. Van Nieuwerburgh. 2018. A Macroeconomic Model with Financially Constrained Producers and Intermediaries. NBER Working Paper No. 24757.

Engel, C. 2016. Exchange rates, interest rates, and the risk premium. *American Economic Review* 106:436–74.

Fama, E. F. 1984. Forward and spot exchange rates. *Journal of Monetary Economics* 14:319–38.

Fama, E. F., and R. R. Bliss. 1987. The information in long-maturity forward rates. *The American Economic Review* 680–92.

Fama, E. F., and J. D. MacBeth. 1973. Risk, Return, and Equilibrium: Empirical Tests. *Journal of Political Economy* 81:607–36.

Fermat Capital. 2015. Insurance-Linked Securities (ILS): " Taller " than you might think.

Friewald, N., R. Jankowitsch, and M. G. Subrahmanyam. 2012. Illiquidity or credit deterioration: A study of liquidity in the US corporate bond market during financial crises. *Journal of Financial Economics* 105:18–36.

Froot, K. A. 1989. New hope for the expectations hypothesis of the term structure of interest rates. *The Journal of Finance* 44:283–305.

———. 2001. The market for catastrophe risk: a clinical examination. *Journal of Financial Economics* 60:529–71.

Froot, K. A., and P. G. O'Connell. 1999. The pricing of US catastrophe reinsurance. In K. A. Froot, ed., *The Financing of Catastrophe Risk*, 195–232. Chicago, IL: University of Chicago Press.

Froot, K. A., and P. G. J. O'Connell. 2008. On the pricing of intermediated risks: Theory and application to catastrophe reinsurance. *Journal of Banking and Finance* 32:69–85.

Fuhrer, J. C. 1996. Monetary policy shifts and long-term interest rates. *The Quarterly Journal of Economics* 111:1183–209.

Gabaix, X., A. Krishnamurthy, and O. Vigneron. 2007. Limits of arbitrage: Theory and evidence from the mortgage-backed securities market. *The Journal of Finance* 62:557–95.

Gârleanu, N., L. H. Pedersen, and A. M. Poteshman. 2009. Demand-based option pricing. *The Review of Financial Studies* 22:4259–99.

Ge, S. 2019. Internal Capital Markets and Product Pricing: Evidence from Weather and Life Insurance Premiums. New York University Working Paper.

Ghysels, E., C. Horan, and E. Moench. 2018. Forecasting through the rear-view mirror: Data revisions and bond return predictability. *Review of Financial Studies* 31:678–714.

Giacoletti, M., K. Laursen, and K. J. Singleton. 2016. Learning, dispersion of beliefs, and risk premiums in an arbitrage-free term structure model. Stanford Working Paper.

Gibbons, M. R., and W. Ferson. 1985. Testing asset pricing models with changing expectations and an unobservable market portfolio. *Journal of Financial Economics* 14:217–36.

Greenwood, R., and D. Vayanos. 2014. Bond supply and excess bond returns. *The Review of Financial Studies* 27:663–713.

Gürkaynak, R. S., B. Sack, and J. H. Wright. 2007. The us treasury yield curve: 1961 to the present. *Journal of Monetary Economics* 54:2291–304.

Haddad, V., and D. A. Sraer. 2019. The Banking View of Bond Risk Premia. NBER Working Paper No. 26369.

Hansen, L. P. 1982. Large sample properties of generalized method of moments estimators. *Econometrica* 50:1029–54.

Hansen, L. P., and R. J. Hodrick. 1980. Forward exchange rates as optimal predictors of future spot rates: An econometric analysis. *Journal of Political Economy* 88:829–53.

———. 1983. Risk averse speculation in the forward foreign exchange market: An econometric analysis of linear models. In J. A. Frenkel, ed., *Exchange Rates and International Macroeconomics*, 113–52. Chicago, IL: University of Chicago Press.

He, Z., B. Kelly, and A. Manela. 2017. Intermediary asset pricing : New evidence from many asset classes. *Journal of Financial Economics* 126:1–35.

He, Z., and A. Krishnamurthy. 2012. A Model of Capital and Crises. *The Review of Economic Studies* 79:735–77.

———. 2013. Intermediary Asset Pricing. *American Economic Review* 103:732–70.

Hong, H., F. W. Li, and J. Xu. 2019. Climate risks and market efficiency. *Journal of Econometrics* 208:265–81.

Ibragimov, R., D. Jaffee, and J. Walden. 2009. Nondiversification traps in catastrophe insurance markets. *The Review of Financial Studies* 22:959–93.

Ilhan, E., P. Krueger, Z. Sautner, and L. T. Starks. 2020. Institutional Investors' Views and Preferences on Climate Risk Disclosure. Swiss Finance Institute Research Paper No. 19-66.

Inoue, A., and L. Kilian. 2005. In-sample or out-of-sample tests of predictability: Which one should we use? *Econometric Reviews* 23:371–402.

Jagannathan, R., and Z. Wang. 1996. The Conditional CAPM and the Cross-Section of Expected Returns. *The Journal of Finance* 51:3–53.

Jiang, Z., A. Krishnamurthy, and H. Lustig. 2018. Foreign safe asset demand and the dollar exchange rate. NBER Working Paper No. 24439.

Jordà, Ò. 2005. Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review* 95:161–82.

Jordà, Ò., M. Schularick, and A. M. Taylor. 2017. Macrofinancial History and the New Business Cycle Facts. In M. Eichenbaum and J. A. Parker, eds., *NBER Macroeconomics Annual*, vol. 31, 213–63. Chicago, IL: University of Chicago Press.

Joslin, S., M. Priebsch, and K. J. Singleton. 2014. Risk premiums in dynamic term structure models with unspanned macro risks. *The Journal of Finance* 69:1197–233.

Jotikasthira, C., A. Le, and C. Lundblad. 2015. Why do term structures in different currencies co-move? *Journal of Financial Economics* 115:58–83.

Kessler, S., and B. Scherer. 2009. Varying risk premia in international bond markets. *Journal of Banking & Finance* 33:1361–75.

Knoll, K., M. Schularick, and T. Steger. 2017. No price like home: Global house prices, 1870-2012. *American Economic Review* 107:331–53.

Koijen, R. S., H. Lustig, and S. Van Nieuwerburgh. 2017. The cross-section and time series of stock and bond returns. *Journal of Monetary Economics* 88:50–69.

Koijen, R. S. J., and M. Yogo. 2015. The cost of financial frictions for life insurers. *American Economic Review* 105:445–75.

———. 2016. Shadow Insurance. *Econometrica* 84:1265–87.

———. 2018. The Fragility of Market Risk Insurance. NBER Working Paper No. 24182.

Krishnamurthy, A., and A. Vissing-Jorgensen. 2011. The effects of quantitative easing on long-term interest rates. *Brookings Papers on Economic Activity* 2:215–65.

———. 2012. The Aggregate Demand for Treasury Debt. *Journal of Political Economy* 120:233–67.

Krueger, P., Z. Sautner, and L. T. Starks. 2019. The Importance of Climate Risks for Institutional Investors. Swiss Finance Institute Research Paper No. 18-58.

Lettau, M., and S. Ludvigson. 2001. Resurrecting the (C)CAPM : A Cross-Sectional Test When Risk Premia Are Time-Varying. *Journal of Political Economy* 109:1238–87.

Lewellen, J., S. Nagel, and J. Shanken. 2010. A skeptical appraisal of asset pricing tests. *Journal of Financial Economics* 96:175–94.

Litterman, R., and J. Scheinkman. 1991. Common factors affecting bond returns. *Journal of Fixed Income* 1:54–61.

Lustig, H., N. Roussanov, and A. Verdelhan. 2011. Common risk factors in currency markets. *The Review of Financial Studies* 24:3731–77.

Lustig, H., A. Stathopoulos, and A. Verdelhan. 2017. The term structure of currency carry trade risk premia. Stanford University Working Paper.

Nagel, S. 2016a. The liquidity premium of near-money assets. *The Quarterly Journal of Economics* 131:1927–71.

———. 2016b. The liquidity premium on near-money assets. *The Quarterly Journal of Economics* 131:1927–71.

Nelson, C. R., and A. F. Siegel. 1987. Parsimonious modeling of yield curves. *Journal of Business* 473–89.

Newey, W. K., and K. D. West. 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55:703–8.

Pancost, N. A. 2018. Zero-coupon yields and the cross-section of bond prices. University of Texas at Austin McCombs School of Business Working Paper.

Pericoli, M., and M. Taboga. 2012. Bond risk premia, macroeconomic fundamentals and the exchange rate. *International Review of Economics & Finance* 22:42–65.

Piazzesi, M., J. Salomao, and M. Schneider. 2015. Trend and cycle in bond premia. Stanford University Working Paper.

Rime, D., A. Schrimpf, and O. Syrstad. 2017. Segmented money markets and covered interest parity arbitrage. Norges Bank Research Paper 15-2017.

RMS. 2012. Cat Bonds Demystified - RMS Guide to the Asset Class.

Roll, R. 1977. A critique of the asset pricing theory's tests. *Journal of Financial Economics* 4:129–76.

Santos, T., and P. Veronesi. 2018. Leverage. Columbia Business School Working Paper.

Sarno, L., P. Schneider, and C. Wagner. 2012. Properties of foreign exchange risk premiums. *Journal of Financial Economics* 105:279–310.

———. 2016. The economic value of predicting bond risk premia. *Journal of Empirical Finance* 37:247–67.

Sekkel, R. 2011. International evidence on bond risk premia. *Journal of Banking & Finance* 35:174–81.

Sen, I. 2019. Regulatory Limits to Risk Management. Harvard Business School Working Paper.

Siriwardane, E. N. 2019. Limited Investment Capital and Credit Spreads. *The Journal of Finance* 74:2303–47.

Svensson, L. E. O. 1995. Estimating forward interest rates with the extended nelson and siegel method. *Sveriges Riksbank Quarterly Review* 1995:13–26.

Swiss Re. 2012. Insurance-linked securities - market update.

Taylor, J. B. 1993. Discretion versus policy rules in practice. In *Carnegie-Rochester Conference Series on Public Policy*, vol. 39, 195–214. Elsevier.

Thornton, D. L., and G. Valente. 2012. Out-of-sample predictions of bond excess returns and forward rates: An asset allocation perspective. *The Review of Financial Studies* 25:3141–68.

Valchev, R. 2017. Bond convenience yields and exchange rate dynamics. Boston College Working Paper.

van Binsbergen, J. H., W. F. Diamond, and M. Grotteria. 2019. Risk-Free Interest Rates. NBER Working Paper No. 26138.

Warga, A., and I. Welch. 1993. Bondholder Losses in Leveraged Buyouts. *The Review of Financial Studies* 6:959–82.

Welch, I., and A. Goyal. 2007. A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies* 21:1455–508.

Wright, J. H. 2011. Term premia and inflation uncertainty: Empirical evidence from an international panel dataset. *American Economic Review* 101:1514–34.

Zhu, X. 2015. Out-of-sample bond risk premium predictions: A global common factor. *Journal of International Money and Finance* 51:155–73.