

Failure to Share Natural Disaster Risk

Tuomas Tomunen*

November 13, 2019

Job market paper

[[Latest version at tuomastomunen.com](http://tuomastomunen.com)]

Abstract

I test whether asset prices reflect the preferences 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 suggesting 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.

Keywords: Risk Sharing, Intermediary Asset Pricing, Reinsurance, Catastrophe Risk, Securitization

JEL Classification: G12, G22

*Tuomas Tomunen is a PhD candidate in Finance at Columbia Business School. tuomas.tomunen@columbia.edu. I thank Kent Daniel, Matthieu Gomez, Franz Hinzen, Bob Hodrick, Harrison Hong, Petri Jylhä, Lira Mota, Pradeep Muthukrishnan, Stijn Van Nieuwerburgh, Tano Santos, Suresh Sundaresan, Paul Tetlock, and Laura Veldkamp for helpful comments. All errors are my own.

1 Introduction

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 risks¹ and are mainly held by specialized asset managers 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 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 well-diversified outside investors are marginal—regardless of

¹Cat 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.

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.




Figure 1: **Observed vs. predicted risk premium.** 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.

Interpreting the pattern in Figure 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 not systematically affect consumption or GDP

²A frictionless macro-finance model predicts that the observations in Figure 1 line up on the x-axis.

growth, house prices, or stock market returns. Furthermore, in a shorter sample starting in 2005, I show that realized returns on cat bonds are not significantly correlated with 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 only specialize in this asset class³, 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 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

³Aon 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.

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

⁴See 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.

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

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. Furthermore, 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 and have more standardized deal characteristics instead of bilateral contracts between the buyer and the seller of protection. This 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.⁶

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.⁷ 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 2 discusses the institutional background of catastrophe bond markets and describes the data. Section 3 develops a simple intermediary asset pricing model and derives several pricing predictions. Section 4 develops empirical procedures to test these predictions. Section 5 provides the main results and evidence on the plausibility of the assumptions needed to rule out other risk-based stories. Section 6 discusses

⁶Other papers discussing alternative risk transfer mechanisms in catastrophe risk space Cummins, Lalonde, and Phillips (2004), Ibragimov, Jaffee, and Walden (2009), and Braun (2016).

⁷See also Daniel, Litterman, and Wagner (2018), and Hong, Li, and Xu (2019), Ilhan et al. (2019), and Krueger, Sautner, and Starks (2019).

and tests various alternative hypotheses for the main results. Section 7 concludes.

2 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.

2.1 Catastrophe bond markets

2.1.1 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 A1 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 A2 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.

2.1.2 Primary market structure

Figure 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 A3 plots the number of issuances in my full sample with different number of tranches.

2.1.3 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 A4 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.

2.1.4 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. Fur-

thermore, 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.

2.1.5 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.⁸

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 billion⁹—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 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 sector, Figure 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

⁸Aon Benfield (2018).

⁹As of October 2019. Source: artemis.bm Insurance Linked Securities Investment Managers & Funds Directory.

them seems a good, abstracted description of the market.

2.2 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.¹⁰ All the matches are then manually checked. Finally, ISIN numbers are manually searched from 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.

¹⁰CUSIP 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.

Secondary market transactions are obtained from TRACE and are available since 2002.¹¹ 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 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.

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

¹¹Even 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.

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.

3 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).

3.1 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

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

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

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

$$X_{i,1} = P_{i,1}(x_i) + C_i, \quad (2)$$

and the expected excess return is

$$E_0(R_i^e) = s_i - el_i. \quad (3)$$

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.

3.2 Investors

3.2.1 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 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 5.3).

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):

$$M_1 \perp x_i \forall i. \quad (4)$$

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 6.2 to rule out such alternative explanation.

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

$$E_0(R_i^e) = 0. \quad (5)$$

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. (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 5.4, 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 5.3 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.

3.2.2 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 1¹², 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.¹³ 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):

$$w_{m,0} \geq \alpha w_0 . \tag{6}$$

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:

$$U(w_{m,1}) = E_0(w_{m,1}) - \frac{\rho}{2} Var(w_{m,1}) . \tag{7}$$

¹²Note 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.

¹³In 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.

3.3 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

$$w_1 = \sum_i q_i (P_{i,1}(x_i) + C_i) + \left(\frac{w_{m,0}}{\alpha} - \sum_i q_i \right) (1 + r). \quad (8)$$

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:

$$E_0(R_i^e) = \beta_i E_0(R_{cat}^e), \quad (9)$$

where $E_0(R_{cat}^e)$ is the expected excess return on the cat bond market portfolio.¹⁴ 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:

$$E_0(R_{cat}^e) = \alpha a \rho \text{Var}_0(R_{cat}^e), \quad (10)$$

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.

¹⁴See Appendix B for proof.

If, on the other hand, the intermediaries are constrained and $\alpha > 0$, the cat bond market specific risk has a positive price.

4 Empirical implementation

In this section, I develop an empirical approach to test the cross-sectional prediction given by Eq. (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. (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.

4.1 Expected excess returns

The expression for expected excess return is given by Eq. (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.

4.1.1 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.

4.1.2 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 $el_{i,new} = el_i + (\bar{x}_{i,new} - \bar{x}_i)$ and $\underline{x}_{i,new} = \max(0, \underline{x}_i + (\bar{x}_{i,new} - \bar{x}_i))$.

4.1.3 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.

4.1.4 Measuring risk-free rate

One additional assumption in Eq. (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, 2016). 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.

4.2 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. (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.¹⁵ 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.

¹⁵See Froot and O'Connell (1999) for somewhat similar approach to estimate catastrophe losses using Monte Carlo simulation.

4.2.1 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 the simulated realized excess return on bond i as:

$$R_{i,t+1}^e(x_i) = (1 + ref_{i,t} + dm_{i,t})P_i(x_i) - r_t, \quad (11)$$

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 :

$$P_i(x_i) = \begin{cases} 1, & x_i > \bar{x}_i \\ \left(\frac{x_i - \underline{x}_i}{\bar{x}_i - \underline{x}_i}\right)^{\phi_i}, & \underline{x}_i \leq x_i \leq \bar{x}_i \\ 0, & x_i < \underline{x}_i \end{cases}, \quad (12)$$

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 A5 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 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.

4.2.2 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.

4.2.3 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 2.1.5). 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.

4.2.4 Simulation results

For each time period and peril category, I simulate 500,000 peril realizations, and calculate bond returns using Eq. (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 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 A1 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 A1 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 A1 are related to pricing results that will be discussed in the next section.

5 Empirical evidence on the failure of risk sharing

In this section, I provide my main cross-sectional results that test Eq. (9). Furthermore, I discuss aggregate catastrophe risk and its time series properties in the light of Eq. (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.

5.1 Main cross-sectional results

Table 3 reports the results of cross-sectional regressions of the following form:

$$E_t(R_{i,t+1}^e) = \lambda_{0,t} + \lambda_{cat,t} \hat{\beta}_{i,t} + \varepsilon_{i,t}, \quad (13)$$

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 $\lambda_{cat,t} - E_t(R_{cat,t+1}^e)$ 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 $E_t(R_{cat,t+1}^e)$. 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.¹⁶

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 4.2, 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 $\lambda_{cat,t} = E_t(R_{cat,t+1}^e)$ easier to reject. Finally, the table reports the cross-sectional R^2

¹⁶As 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.

values for each year. They range between 0.29 and 0.84, suggesting a reasonably high fit.¹⁷

Table A1 illustrates how the measurement error affects the results. As discussed in Section 4.2.4, 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 A6 repeats the analysis presented in Figure 1 and highlights these five example observations, showing clearly the outlier bond.

As discussed earlier in 4.2.4, 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 A6 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.

5.2 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 are

$$U(w_m) = \frac{w_m^{1-\hat{\rho}} - 1}{1 - \hat{\rho}}. \quad (14)$$

¹⁷Table A2 repeats the analysis with actual trade prices from TRACE instead of the sheet prices. The results are consistent with those in Table 3, with $\lambda_{cat,t} = 0$ being rejected for most of the years. Because of lack of trading 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.

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. (10) for this expression and rearranging gives:

$$\frac{E_t(R_{cat,t+1}^e)}{Var_t(R_{cat,t+1}^e)} = \hat{\rho} \frac{Size_t}{AUM_t}, \quad (15)$$

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. (15) using the simulation framework introduced in Section 4.2. This gives me the observed premium per unit of risk. The right-hand side of Eq. (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 12.2.¹⁸

Figure 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. 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

¹⁸Unfortunately, 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%.

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.

5.3 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 countries¹⁹ 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).²⁰ The database contains event- and country-level information on mass disasters from 1900 to the present.²¹

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

¹⁹Australia, Belgium, Canada, France, Italy, Japan, Netherlands, Norway, Portugal, Spain, Switzerland, United Kingdom, United States.

²⁰EM-DAT: The Emergency Events Database - Université catholique de Louvain (UCL) - CRED, D. Guha-Sapir - www.emdat.be, Brussels, Belgium.

²¹Since 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.

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

$$dmg_{i,t} = \frac{\sum_d DMG_{d,i,t}}{GDP_{i,t-1}}. \quad (16)$$

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

$$\Delta_h y_{i,t+h} = \gamma_i + \gamma_t + b_1 Small_{i,t} + b_2 Large_{i,t} + \varepsilon_{i,t}, \quad (17)$$

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, 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 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 are 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 A3 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.²² 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.

5.4 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

²²Note 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.

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 and Table 3. In order to explain the pattern, the estimates would need to be more downward biased for bonds with higher betas.

To evaluate the accuracy of the actuarial models, I estimate the loss distributions on cat bonds using the simulation framework introduced in Section 4.2 for each year in the sample (2003-2018). Then, I calculate the cumulative loss distribution on a strategy that invests \$1,000 on an equally weighted portfolio of single-peril cat bonds rebalanced annually at the end of June.²³ 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 6 plots the expected cumulative dollar losses implied by the actuarial models.²⁴ 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 models are unbiased, although the confidence intervals are relatively large due to the highly skewed distributions of the events. For example, even if the

²³Because 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.

²⁴Figure A7 shows the results as a percentage of capital invested.

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.²⁵ To sum up, while we 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 3.2.1: 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. (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 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).

6 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.

²⁵Assuming market conditions stay similar to 2018.

6.1 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.²⁶ 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.

6.1.1 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 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.²⁷ Somewhat more crude measures, based on bond characteristics, include issuance size, bond age, and time to maturity.

Table 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

²⁶He 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.

²⁷Number 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.

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 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.

6.1.2 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 index²⁸ 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 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

²⁸Eurekahedge ILS Advisers Index.

should be willing to allocate capital to such a fund without imposing high requirements on the expected rate of return.

6.1.3 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,²⁹ 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 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 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

²⁹Bond triggers if the underwriting losses of the issuing (re)insurance company are above some threshold, see Section 2.1.4 for description of the trigger types.

the results, implying a rejection of the hypothesis that that moral hazard or adverse selection is a key determinant of the pricing results.

6.1.4 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.

6.2 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 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.³⁰ 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:

$$\Delta s_{i,j} = \lambda \Delta el_{i,j} + \varepsilon_{i,j}, \quad (18)$$

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.³¹ 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, 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.

³⁰Note 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.

³¹The results are robust to clustering standard errors at month-level, although this specification amplifies the potential small-sample issues for most restricted specifications.

Table 8 shows the formal results for an OLS estimation of Eq. (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.³² 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.

6.3 Other explanations

Table 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 6.1.3). 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.

³²While 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.

7 Conclusion

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.

References

- Acharya, V. V., T. Philippon, and M. Richardson. 2016. Measuring Systemic Risk for Insurance Companies. In F. Hufeld, R. S. J. Kojien, 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.
- Aon Benfield. 2009. Insurance-Linked Securities - Adapting to an Evolving Market.
- . 2018. Insurance-Linked Securities - Alternative Capital Fortifies Its Position.
- Bansal, R., D. Kiku, and M. Ochoa. 2017. Price of Long-Run Temperature Shifts in Capital Markets. Working Paper, Duke University.
- Bao, J., J. Pan, and J. Wang. 2011. The Illiquidity of Corporate Bonds. *The Journal of Finance* 66:911–46.
- Barro, R. J. 2006. Rare Disasters and Asset Markets in the Twentieth Century. *The Quarterly Journal of Economics* 121:823–66.
- Bekaert, G., R. J. Hodrick, and D. A. Marshall. 2001. Peso problem explanations for term structure anomalies. *Journal of Monetary Economics* 48:241–70.
- Bolton, P., and M. Kacperczyk. 2019. Do Investors Care about Carbon Risk? Working Paper, Columbia University.
- Braun, A. 2016. Pricing in the Primary Market for Cat Bonds: New Empirical Evidence. *Journal of Risk and Insurance* 83:811–47.
- 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.
- 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.
- 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.

- Daniel, K. D., R. B. Litterman, and G. Wagner. 2018. Applying Asset Pricing Theory to Calibrate the Price of Climate Risk. Working Paper, Columbia Business School.
- Dick-Nielsen, J. 2009. Liquidity Biases in TRACE. *The Journal of Fixed Income* 19:43–55.
- . 2014. How to clean Enhanced TRACE data. Working Paper, Copenhagen Business School.
- 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.
- Elenev, V., T. Landvoigt, and S. Van Nieuwerburgh. 2018. A Macroeconomic Model with Financially Constrained Producers and Intermediaries. Working Paper, NBER.
- 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. 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.
- 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. Working Paper, New York University.

- Greenwood, R., and D. Vayanos. 2014. Bond supply and excess bond returns. *The Review of Financial Studies* 27:663–713.
- Haddad, V., and D. A. Sraer. 2019. The Banking View of Bond Risk Premia. Working Paper, NBER.
- 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. 2019. Institutional Investors’ Views and Preferences on Climate Risk Disclosure. Working Paper, Frankfurt School of Finance & Management.
- Jagannathan, R., and Z. Wang. 1996. The Conditional CAPM and the Cross-Section of Expected Returns. *The Journal of Finance* 51:3–53.
- 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.
- 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. 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. Working Paper, New York University.

- Krishnamurthy, A., and A. Vissing-Jorgensen. 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. Working Paper, University of Geneva.
- 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.
- Lustig, H., N. Roussanov, and A. Verdelhan. 2011. Common risk factors in currency markets. *The Review of Financial Studies* 24:3731–77.
- Nagel, S. 2016. The liquidity premium on near-money assets. *The Quarterly Journal of Economics* 131:1927–71.
- 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. Working Paper, NBER.
- Sen, I. 2019. Regulatory Limits to Risk Management. Working Paper, London Business School.
- Siriwardane, E. N. 2019. Limited Investment Capital and Credit Spreads. *The Journal of Finance* 74:2303–47.
- Swiss Re. 2012. Insurance-linked securities - market update.
- van Binsbergen, J. H., W. F. Diamond, and M. Grotteria. 2019. Risk-Free Interest Rates. Working Paper, University of Pennsylvania.
- Warga, A., and I. Welch. 1993. Bondholder Losses in Leveraged Buyouts. *The Review of Financial Studies* 6:959–82.

Tables and figures

Table 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 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 3: Pricing of catastrophe market risk

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

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

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 $\lambda_{cat,t} - E_t(R_{cat,t+1}^e)$ 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)	$\lambda_{cat,t} - E_t(R_{cat,t+1}^e)$	(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 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 5: Pricing of catastrophe market risk and liquidity

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

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

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*** (0.15)	1.24*** (0.15)	1.36*** (0.14)	1.29*** (0.14)	1.40*** (0.15)	1.39*** (0.15)
λ_{cat}	2.03*** (0.20)	2.06*** (0.20)	1.81*** (0.17)	2.08*** (0.20)	1.99*** (0.20)	2.00*** (0.20)
λ_{liq}	-0.20** (0.09)	-0.37*** (0.10)	-0.10 (0.10)	-0.30*** (0.08)	-0.15*** (0.05)	0.03 (0.05)
N	57	57	57	57	57	57
R^2	0.54	0.54	0.53	0.55	0.54	0.54

Table 6: Explaining realized cat bond manager returns with other asset classes

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

$$R_{ILS,t}^e = a_i + b_i R_{i,t}^e + \varepsilon_{i,t},$$

where the dependent variable is monthly excess return on Eureka-hedge 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** (0.136)	0.258* (0.139)	0.266* (0.158)	0.319** (0.131)	0.268* (0.160)
Equities	0.027 (0.021)				0.005 (0.034)
High-yield bonds		0.057* (0.031)			0.038 (0.051)
MBS			0.091 (0.173)		0.139 (0.165)
Carry trade				0.031 (0.040)	0.008 (0.049)
N	156	156	156	154	154
R^2	0.013	0.026	0.005	0.006	0.034

Table 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

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

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 $\lambda_{cat,t} - E_t(R_{cat,t+1}^e)$ 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 3. "Noncallables" and "Earthquake" columns include only noncallable 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*** (0.13)	1.20*** (0.13)	1.50*** (0.18)	1.28*** (0.15)
$\lambda_{cat,t}$	2.06*** (0.18)	2.11*** (0.18)	1.98*** (0.30)	1.94*** (0.25)
$\lambda_{cat,t} - E_t(R_{cat,t+1}^e)$	-1.10*** (0.12)	-1.05*** (0.12)	-1.18*** (0.30)	-1.23*** (0.19)
N	63	63	63	63
R^2	0.49	0.51	0.43	0.49

Table 8: Price of non-extreme catastrophe 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 . $\text{Min}(el_A)$ denotes the sample selection criteria for a given regression. For example, $\text{Min}(el_A)=1.00$ implies that only those pairs of cat bond tranches where the safer tranche 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 ^{***} (0.09)	1.42 ^{***} (0.09)	1.48 ^{***} (0.12)	1.44 ^{***} (0.13)	1.44 ^{***} (0.15)	1.49 ^{**} (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
$\text{Min}(el_A)$	0.00	0.50	1.00	1.50	2.00	2.50

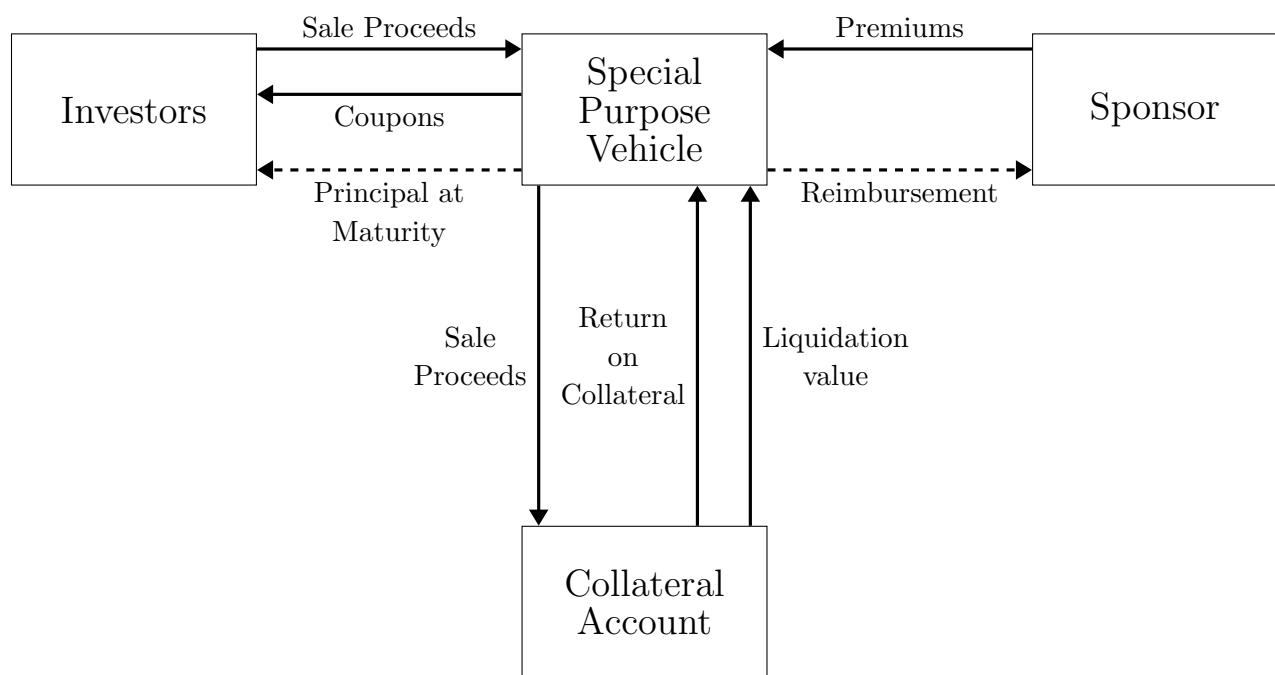


Figure 2: Typical cat bond deal structure

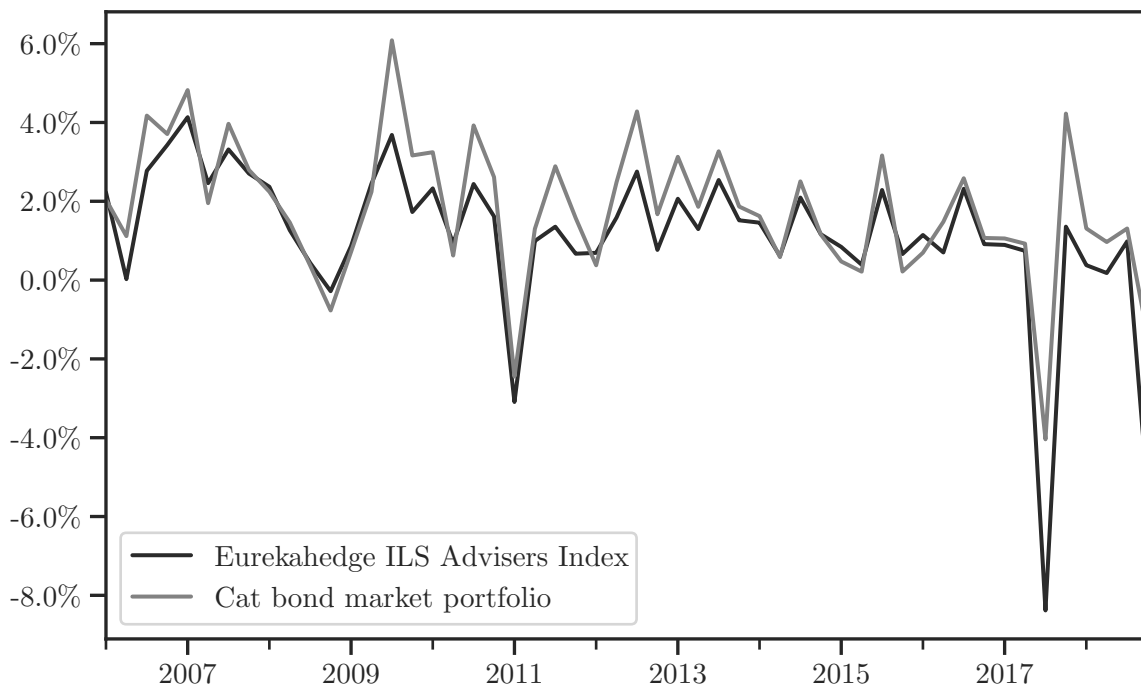


Figure 3: Returns on specialist funds and value-weighted portfolio of sample bonds




Figure 4: **Time series evolution of price of natural disaster risk.** The figure plots observed and predicted premium on natural disaster risk, where coefficient of relative risk aversion $\hat{\rho}=6.1$. R_{cat}^e 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.


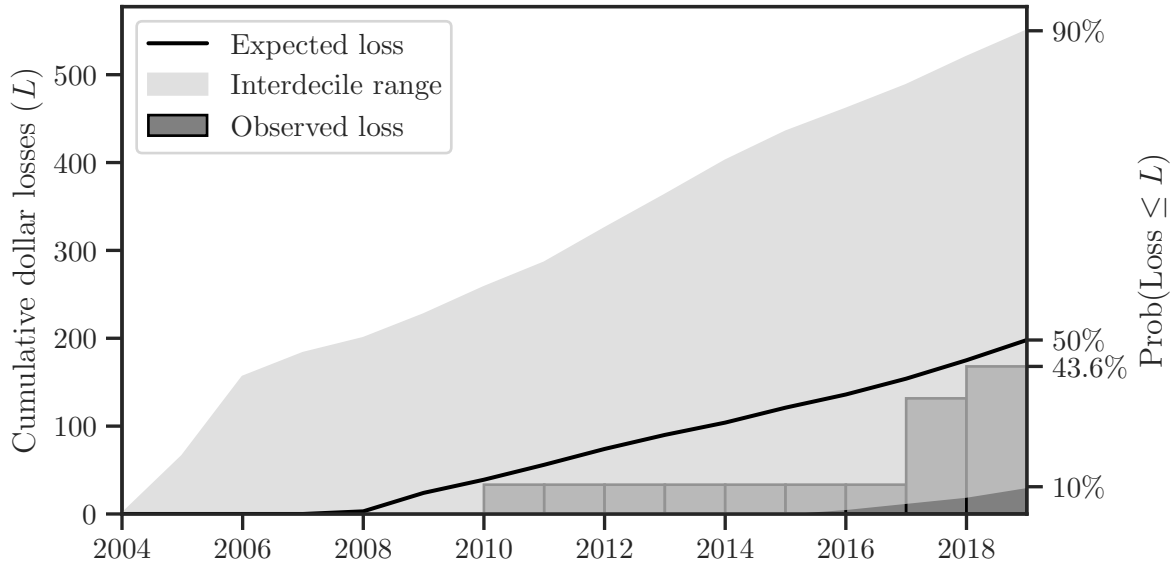



Figure 5: **Impulse-responses of selected macroeconomic variables to natural disasters occurring in year 0.** Shaded region shows 95% confidence interval on the estimates. Standard errors are clustered by year. Sample includes 13 developed countries from 1950 to 2016.



(a) Unbiased actuarial models



(b) Biased actuarial models

Figure 6: **Predicted and actual losses on cat bonds due to natural disasters.** 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 7: Differences in spreads and expected losses between two cat bond tranches

Appendix A: Data

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.³³ 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) 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.

³³In 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).

Appendix B: Proof of main predictions

Manager f 's optimization problem can be written as:

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

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:

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

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

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

where $P_{cat,1} = \mathbf{P}_1\theta$. Note that by definition, $P_{i,0} = 1 \forall i$. As a result,

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

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,

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

and

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

Note that since the $E_0(R_{cat}^e) > 0$, constraint in Eq. (6) is always binding. This is because $E_0(M_1 R_{cat}^e) > 0$ and as a result the outside investors are willing to give the managers the maximum amount of capital possible.

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 $M(x_i, \mathbf{z})$, 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. (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 .³⁴

Now, let us make two assumptions.

Assumption A1 (Small disasters are not priced):

$$M(x_i, \mathbf{z}) = M(\mathbf{z}), x_i \geq x_i^*. \quad (19)$$

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

$$\underline{x}_A \geq x_A^*. \quad (20)$$

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.

³⁴In 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.

Figure A8 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 3.2.1 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:

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

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

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

Now, since $y_i = r + s_i$ and $E_0 (X_i) = (1 - el_i)$, 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 $\Delta el = \ln(1 - el_A) - \ln(1 - el_B)$.

Appendix: Tables and figures

Table A1: 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 . $\hat{\beta}_{i,t}E_t(R_{cat}^e)$ is predicted premia, and $E_t(R_i^e)$ observed premia. Premium for cat bond market portfolio $E_t(R_{cat}^e)$ 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	$\hat{\beta}_{i,t}$	$\hat{\beta}_{i,t}E_t(R_{cat}^e)$	$E_t(R_i^e)$
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 A2: Pricing of catastrophe market risk

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

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

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 $\lambda_{cat,t} - E_t(R_{cat,t+1}^e)$ 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)	$\lambda_{cat,t} - E_t(R_{cat,t+1}^e)$	(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 A3: Effect of natural disasters on selected macroeconomic variables

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

$$\Delta_h y_{i,t+h} = \gamma_i + \gamma_t + b_1 \text{Small}_{i,t} + b_2 \text{Large}_{i,t} + \varepsilon_{i,t},$$

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. $\text{Small}_{i,t}$ 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 . $\text{Large}_{i,t}$ 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			
	$h = 0$		$h = 2$		$h = 0$		$h = 2$	
$\text{Small}_{i,t}$	-0.011 (0.227)	-0.259 (0.258)	0.088 (0.616)	-1.067 (0.718)	0.198 (0.234)	0.056 (0.298)	0.526 (0.625)	-0.758 (0.775)
$\text{Large}_{i,t}$	0.820 (0.518)	0.896 (0.601)	1.917* (1.136)	2.598* (1.565)	0.522 (0.583)	0.950 (0.738)	1.639 (1.281)	2.754 (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			
	$h = 0$		$h = 2$		$h = 0$		$h = 2$	
$\text{Small}_{i,t}$	-0.161 (0.921)	-0.395 (0.916)	2.086 (2.532)	-0.529 (2.667)	0.215 (2.320)	0.571 (2.873)	4.126 (4.547)	5.020 (5.551)
$\text{Large}_{i,t}$	2.444 (2.250)	3.093 (1.971)	3.180 (4.319)	1.834 (4.486)	0.275 (5.155)	-3.811 (5.744)	6.921 (8.239)	7.875 (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 A1: Total face value outstanding for sample bonds




Figure A2: Ten most covered perils




Figure A3: Distribution of number of tranches per issuance

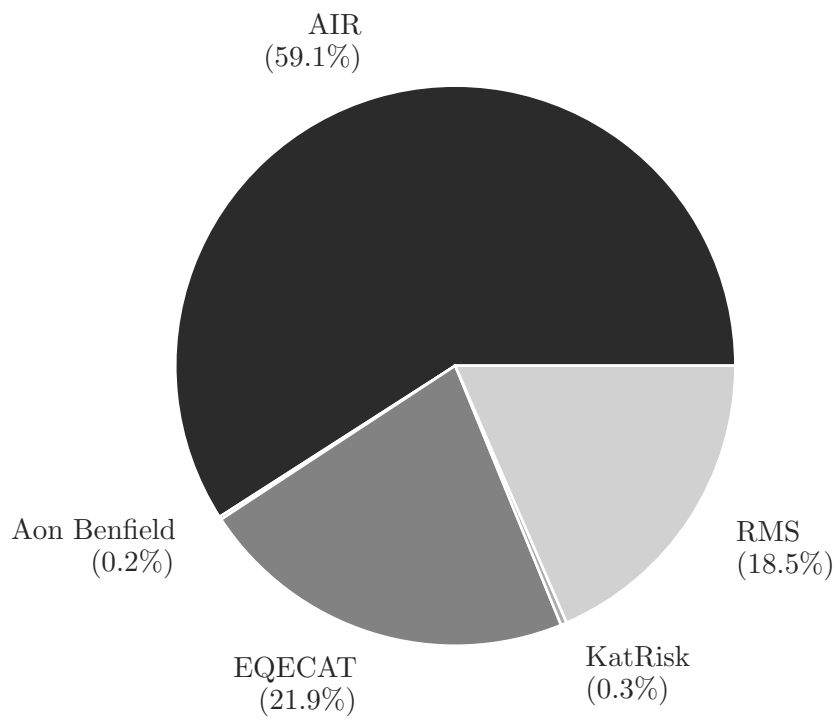


Figure A4: Market share of risk modeling companies




Figure A5: Illustration of loss function P


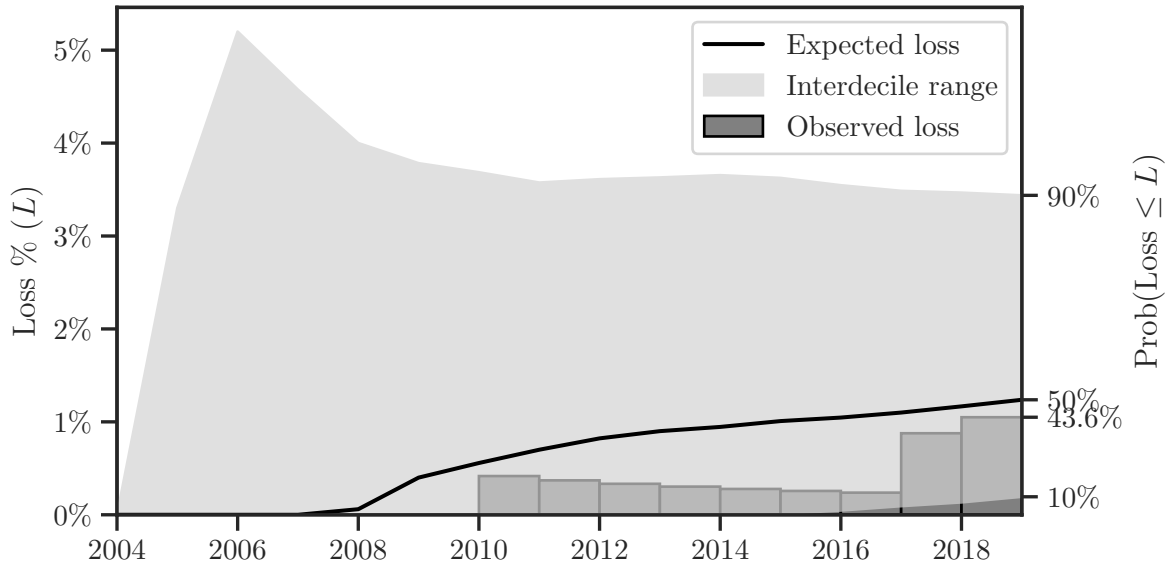



Figure A6: **Observed vs. predicted risk premium.** The figure plots expected excess returns (discount margin - expected actuarial loss) on sample cat bonds between 2003 and 2018 against their predicted values from a single-factor intermediary model. Bonds' betas are estimated from 500,000 years of simulated 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 A1. 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 A7: **Predicted and actual losses on cat bonds due to natural disasters.** 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 A8: Illustration of Assumptions A1 and A2