Tail Overweighting in Catastrophe Bond Pricing

Since the late 20th century, (re)insurers have transferred severe natural catastrophe risks to the capital market via catastrophe bonds. These bonds offer high spreads and significant diversification potential for investors, as the underlying risks are uncorrelated with traditional capital market risks. Explaining the spreads of these bonds remains challenging. Traditional asset pricing models provide limited insights due to the uncorrelated nature. This study introduces a new perspective based on an overweighting of tail components in the underlying loss process. The model is calibrated using unique cyber loss dataset from a leading risk modeling agency. The calibration reveals that the relative tail overweighting in catastrophe bonds increases as the likelihood of a trigger event decreases. The more "pure" tail risk within the bond, the more overemphasized the tail components in the pricing process – similar to the volatility smile observed in options markets. Moreover, combining beta-risk factors with the new tail overweighting approach can provide an almost perfect explanation of catastrophe bond spreads.

Keywords: NatCat Risk, Cyber Risk, Cat Bond, Cyber Bond, Jump Processes, Tail Risk

JEL Classification: D81, G12, G22, Q54

The author declares that he has received no research support and has no financial relationships or other potential conflicts of interest to disclose.

1 Introduction

The increasing frequency and severity of natural disasters, driven by climate change, have expanded the scope of private risk-sharing beyond traditional property and casualty insurance. In addition, the rise of cyber threats due to global networking and digitalization has further broadened this scope. Since the late 20th century, insurance-linked securities, such as catastrophe bonds (cat bonds), have integrated capital markets into the risk-sharing process.[1] These bonds allow (re)insurers or brokers to initially underwrite higher loss layers from policyholders and then transfer those layers to capital markets, effectively becoming bondholders seeking protection in case of a triggering event. Since these risks are largely uncorrelated with capital market risks, they are attractive investments due to their high spreads and diversification potential (Cummins and Weiss, 2009). However, understanding and replicating the spreads of cat bonds remains an unsolved puzzle. Traditional asset pricing models explain only a small part of the risk premium (e.g., Braun et al., 2019). The limitations of traditional beta models can be explained with the findings of Barberis and Huang (2008), who show that even idiosyncratic skewness unrelated to the overall market returns will be overpriced relative to what would be expected in a traditional beta-risk-based economy. Another issue is understanding the underlying risk dynamics. The availability of high-frequency natural catastrophe data is limited, and long-term data is scarce, particularly for newer markets such as cyber. For events like pandemics, data is virtually nonexistent. Moreover, the cat bond market is highly heterogeneous, making it difficult to find bonds with comparable characteristics.

Traditional models struggle to explain cat bond spreads, creating challenges for accurate pricing and risk management in the growing risk transfer market. While most research relies on cross-sectional analyses of the entire cat bond market, this paper examines spreads based on the underlying cumulative distribution function (cdf). It is well known that under probability weighting, individuals tend to overestimate low-probability events (Kai-Ineman and Tversky, 1979). Households, in particular, are willing to pay higher premiums for insurance against rare but severe risks (Sydnor, 2010). Building on this insight, this paper accounts for the underlying loss dynamics and calibrate market prices by emphasizing jump-risk components. Using a unique industry dataset and a series of homogeneous cat bonds with cyber risk as the underlying, this study addresses market heterogeneity and the limited availability of long-term data for calibration. The combination of a relatively new and homogeneous risk within the cat bond market and the availability of the unique industry dataset enables, for the first time, such a cdf-based analysis. Prior studies often aggregated risk perils with different cdfs, such as pooling U.S. hurricane data at the national level instead of analyzing it at the state level, to ensure a sufficient number of data points.

Catastrophic risks in insurance and financial markets are commonly modeled using compounded Poisson processes (Insurance: e.g., Bowers et al., 1986; Lee and Yu, 2002; Ma and Ma, 2013; Financial markets: e.g., Bates, 1996; Andersen et al., 2002; Kou, 2002; Eraker, 2004). Initially, a probability measure $\mathbb{P}$ represents the distribution of loss occurrence and,

[1] In this paper, the term "cat bond" is used generically to include bonds with cyber risks. Some papers distinguish between these types, referring to bonds with cyber risks as "cat cyber bonds" or "cyber bonds." This distinction is omitted here to avoid unnecessary repetition.

in an ideal frictionless market, would align with prices. Industry data from a leading risk modeling agency – currently the most accurate approximation of the real-world probability measure – allow for the estimation of $\mathbb{P}$. While no cyber-linked cat bonds have been triggered yet, these loss data provide a robust foundation for calibrating $\mathbb{P}$ as the "true" underlying loss distribution. However, real markets are not frictionless and complete, requiring the definition of a pricing measure $\mathbb{Q}$. This paper introduces a measure reflecting market prices by overweighting tail risks through a transformation of jump components in the compounded Poisson processes. The approach is based on actuarial probability distortion, specifically the Esscher transformation (Esscher, 1932; Gerber and Shiu, 1994). This distortion captures investors' heightened sensitivity to extreme losses (Bollerslev and Todorov, 2011) and their aversion to negative skewness (Pan, 2002). Since markets are incomplete, infinite $\mathbb{Q}$ measures are possible. To find the true one, the transformation parameters are calibrated on observed spreads and compared. The calibration shows a unique pattern between the transformation parameter and the probability of the first loss.

The calibration reveals that, despite market incompleteness and the infinite choices for defining a $\mathbb{Q}$ measure, probability distortion toward tail risks follows a clear structure. As the likelihood of a bond being triggered decreases, investors systematically overweight tail probabilities in the pricing process. This behavior aligns with broader observations in financial markets, where concerns about downside events shape pricing dynamics. Gao et al. (2019) describe this as a reflection of tail risk concerns, driven by investors' subjective ex-ante beliefs. A similar effect can be observed in other markets; for instance, the options market recognizes this phenomenon as a volatility smile (e.g., Derman and Kani, 1994). In the case of cat bonds, the asymmetry in risk perception results in a half-smile, reflecting the one-sided skewness. The $\mathbb{Q}$ measure captures a distinct pricing structure that extends beyond conventional models, providing a unique metric for tail-risk aversion in pricing. Furthermore, the results indicate that approximately 40% of the spread in the cat bond market with the risk peril "cyber" can be explained by this tail overweighting. When incorporating beta factors – combining cross-sectional risk premiums with the cdf-specific tail overweighting – the spread can be replicated with near-perfect accuracy. This highlights the importance of integrating traditional risk premium models with loss-specific pricing dynamics rather than treating them separately. By unifying these approaches, a more comprehensive and consistent framework for cat bond pricing emerges, improving both the understanding of market behavior and risk assessment.

The findings offer a new perspective on cat bond pricing by showing that cross-sectional risk premiums and cdf-specific tail risk distortions should be combined. This integration leads to a more consistent and accurate pricing framework, improving the understanding of how investors price downside risks and refining existing models to better reflect market behavior. The structured overweighting of tail risks enhances risk assessment and capital allocation in the insurance-linked securities market. Explicitly quantifying tail-risk aversion contributes to more efficient pricing mechanisms, benefiting both cedents and investors. Moreover, the observed probability distortions align with patterns in other financial markets, suggesting a broader relevance beyond cat bonds. Investors systematically adjust for downside risks across asset classes, reinforcing the need for pricing models that incorporate probability distortions.

This paper addresses a significant research gap, as previous studies primarily focus on

cross-sectional analyses. For example, Galeotti et al. (2013) compare several basic factor models and demonstrate that a simple model using expected loss as the sole factor already explains a substantial portion of cat bond spreads. Similarly, Braun (2016) and Herrmann and Hibbeln (2021) find that part of the cat bond spreads can be attributed to cedent characteristics and seasonal factors. Moreover, Beer and Braun (2022) highlight that the probability of first loss is a key predictor in the secondary market, while Gürtler et al. (2016) show that risk perceptions can shift following an adverse event, such as Hurricane Katrina. Another study by Tomunen (2022) constructs a cat bond market factor using a cat bond portfolio. However, it has limitations, particularly in assigning low spreads to underrepresented risks, like hurricanes in New York or emerging risks such as cyber, despite market data suggesting higher spreads. While several studies attempt to explain cat bond pricing by linking loss dynamics to risk-neutral measures (e.g., Bakshi and Madan, 2002; Lee and Yu, 2002), the calibration process remains challenging due to the vast number of possible measure specifications. The perspective taken in this paper, which incorporates tail overweighting, is new to the best of the author's knowledge. The findings extend beyond cat bond literature and are relevant to other markets characterized by downside risk and the limitations of traditional factor models. For instance, similar challenges arise in corporate bond markets, where downside risk is a recognized pricing factor (e.g., Bai et al., 2019; Dickerson et al., 2023). The pricing of insurance products also faces related difficulties. Zanjani (2002) suggests that high markups in insurance pricing are largely driven by capital costs, arising from the lack of correlation between insurance risks and broader capital market risks. While efforts have been made to develop factor models tailored to insurers (e.g., Ben-Ammar et al., 2018), a comprehensive model that fully captures these costs on the asset side is still missing. As a result, suboptimal factor models remain widely used (e.g., Cummins and Phillips, 2005; Braun et al., 2023), indicating that high markups may also be partly attributed to tail risk aversion.

This paper is structured as follows: Section 2 introduces the cat bond structure and the underlying model. Section 3 provides an overview of the institutional setting, describes the data, and presents the model calibration, highlighting evidence of tail overweighting. Section 4 offers concluding remarks.

2 Cat bond structure and model

What are cat bonds?

The issuers of cat bonds, known as cedents, are often internationally operating (re)insurers or insurance brokers. First, the cedent underwrites various insurance policies to build its portfolio. For example, consider the case of diverse hurricane insurance policies written across the state of Florida. While the cedent can diversify many smaller losses within the counties, the fat-tail of the loss distribution remains, consisting of low-frequency, high-severity events. If an event triggers the cedent's entire portfolio, it becomes existentially threatening (e.g., Braun et al., 2023). An example of such an event was Hurricane Milton in October 2024. Therefore, the cedent transfers these risks to the capital market via a cat bond. Based on

the cedent's portfolio, modeling agencies design the cat bond. They estimate key parameters such as the bond's attachment quantile (i.e., the probability of triggering), the expected loss, the spread and the notional (cat bond size). These modeling agencies are specialized services exclusively focused on cat risk modeling, as historical data for low-frequency, high-severity events is scarce. Thus, the modeling process extends beyond historical data, incorporating environmental factors such as climate change, expert opinions, and large-scale simulations. Building on the previous example, assume that the cedent is able to retain losses up to $500 million annually but wishes to transfer the next $300 million to the capital market. The cedent issues a bond and invites investors to contribute $300 million to a special-purpose vehicle (SPV). Investors receive detailed information from the modeling agency, and a trigger is defined. In this case, the trigger would activate if the cedent's loss volume during a given period exceeds $500 million. This is an "indemnity" trigger type, currently accounting for over 70% of the market, followed by an "industry loss index" trigger with more than 20% market share (Artemis, 2025b). The trigger type has no influence on the spread (Gürtler et al., 2016). Once the funds are deposited into the SPV, they are invested risk-free in a trust account. When the bond's period ends or the bond is triggered, the SPV is liquidated. The cedent receives compensation for the incurred loss (if any). Investors receive the remaining (or full) amount in the SPV, the risk-free return from the trust account, and the agreed spread from the cedent. Figure 1 illustrates this interaction.

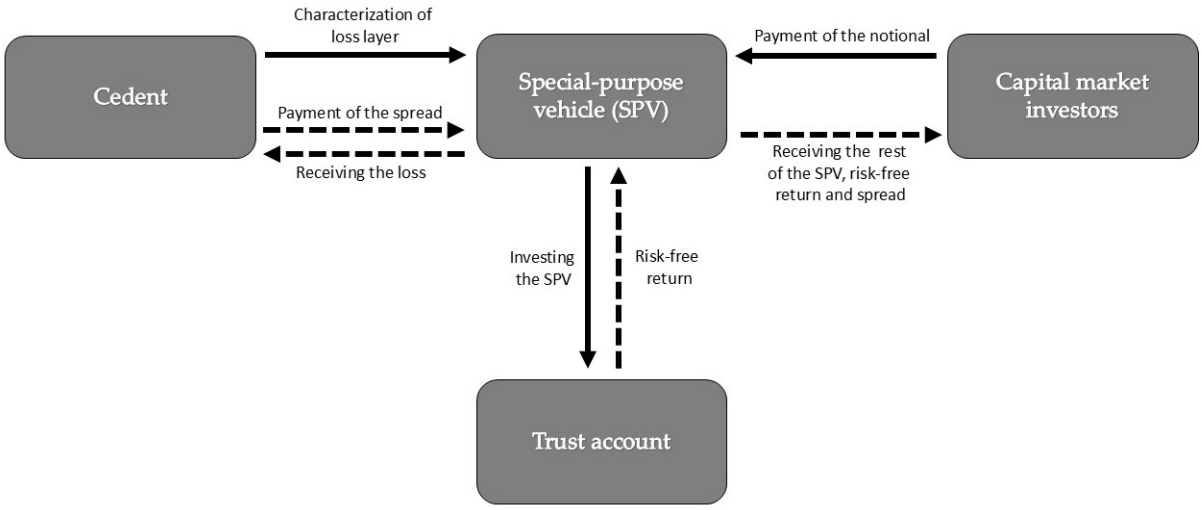


Figure 1: Structure of a cat bond.

This figure illustrates the structure and cash flows of a cat bond. Solid lines represent the processes during bond issuance, while dashed lines indicate the actions when the bond is liquidated; own illustration based on information from Artemis (2025d).

Model

Consider a one-period model where the bond is issued at $t = 0$ and liquidated at $t = 1$. At $t = 1$, the cedent claims a fraction of the loss $\bar{L} \geq 0$, which is deducted from the cash flow to

the investors. This cash flow consists of three components: (1) the initial equity investment by the investors, S, (2) the risk-free interest, $r_f \cdot S$, and (3) a risk premium, $P = s \cdot S$, where s is the spread on the invested equity. While components (1) and (2) are sourced from the SPV, the premium in component (3) is paid by the bond cedent to compensate the investors for their allocated capital.

Let T denote the trigger event and $K = T + S$ represent the cap of the cat bond. The loss claimed by the bond cedent can then be expressed as:

$$\bar{L}_{bond} = \begin{cases} 0 & \text{if } \bar{L} \leq T \\ \bar{L} - T & \text{if } T < \bar{L} \leq K \\ K - T & \text{if } K < \bar{L}, \end{cases}$$

or

$$\bar{L}_{bond} = \max(\bar{L} - T, 0) - \max(\bar{L} - K, 0). \tag{1}$$

From an insurance perspective, cat bonds act as excess-of-loss insurance contracts, where the cedent is analogous to a policyholder, and the capital market functions as the insurer. In this setup, the trigger defines the deductible, the bond size determines the limit (the loss layer covered between the trigger and the cap), and the spread serves as the premium paid by the cedent to transfer the risk. From a different angle, the value of the bond payoff at $t = 0$ can be described using European call options. This aligns with the option pricing framework, as Doherty and Garven (1986) demonstrated that insurance contracts can be priced using European call options.

To replicate the spread of a cat bond, this paper aims to compare bonds with identical loss distribution $\bar{L}$ rather than grouping bonds into broad categories. This is particularly challenging because many cat bonds cover heterogeneous risks, complicating the calibration process. Cat bond losses are primarily situated in the tail of the loss distribution, often encompassing fat-tail behavior. Consequently, incorporating a jump process into the loss dynamics is essential – a standard approach in cat risk modeling (e.g., Lee and Yu, 2002; Ma and Ma, 2013). Let N represent a Poisson process with a constant jump arrival rate λ and jump size Z, where Z takes values $z \in \mathbb{R}$. This process is uncorrelated with the global capital market and specific to the underlying loss, reflecting purely idiosyncratic shocks. The jump sizes are independently and identically distributed as $\mathcal{N}(\alpha, \delta^2)$. The expected proportional unique jump sizes are given by:

$$\kappa_Z = \mathbb{E}_{\mathbb{P}}[\exp(Z) - 1] = \int_{\mathbb{R}} [\exp(Z) - 1] m_{\mathbb{P}}(dz),$$

where $m_{\mathbb{P}}(dz)$ is the density of Z (Merton, 1976). The loss dynamics can then be expressed as:

$$\frac{dL}{L} = (\mu - \lambda\kappa_Z)dt + \sigma dW + \int_{\mathbb{R}} [\exp(z) - 1] p(dz),$$

where $p(\cdot)$ is the Poisson measure, μ is the drift of the loss process and σ is volatility of the loss process. The Poisson measure and the Wiener process are independent. The loss at

$t = 1$ is therefore given by the solution:

$$\bar{L} = L_0 \exp\left(\left(\mu - \lambda\kappa_Z - \frac{\sigma^2}{2}\right) + \sigma W + \sum_{m=1}^{N} Z_m \right), \tag{2}$$

defining the loss process $\mathbb{P}$. Under the assumption of a frictionless and fair market, the spread would be computed under $\mathbb{P}$, given by $s = \frac{\mathbb{E}_{\mathbb{P}}[\bar{L}_{bond}]}{S}$. This formulation implies that the spread reflects the expected loss investors face relative to their initial investment. However, in practice, the spread exceeds the expected loss. For instance, in Q4 2024, the average multiple – indicating how many times the expected loss is paid as the spread – was approximately 2.8 (Artemis, 2025c).

To evaluate pricing conditions, the probability measure $\mathbb{P}$ is transformed into a pricing measure $\mathbb{Q}$, implying that the underlying loss is calibrated under $\mathbb{P}$ but priced under $\mathbb{Q}$. The model here assumes that low-frequency, high-severity losses significantly influence pricing, necessitating the transformation of the jump processes. The measure $\mathbb{Q}$ introduces a probability distortion in the tail, economically representing tail-risk aversion (e.g., Gao et al., 2019). The Esscher transformation proposed by Esscher (1932) is applied, as it is widely used in actuarial and option pricing literature (e.g., Gerber and Shiu, 1994). After the transformation (see Appendix A.1), the change in intensity is defined as:

$$\tilde{\lambda} = \lambda \mathbb{E}_{\mathbb{P}}[\exp(\eta Z)], \tag{3}$$

and the distribution of jump sizes changes, with Z now normally distributed as $\mathcal{N}(\tilde{\alpha}, \delta^2)$, where $\tilde{\alpha} = \alpha + \delta^2\eta$. The transformation parameter η is the fundamental factor in transitioning from $\mathbb{P}$ to $\mathbb{Q}$. In a market with jump components, multiple equivalent martingale measures exist, leading to varied pricing measures (e.g., Harrison and Pliska, 1981). The calibration process in the empirical analysis focuses on η.

Given $\mathbb{Q}$, the spread could now be expressed as $s = \frac{\mathbb{E}_{\mathbb{Q}}[\bar{L}_{bond}]}{S}$. This approach, however, presents a static perspective that appears to diverge from factor models and the traditional beta-risk-based framework commonly applied in financial markets. Nevertheless, this divergence is not the intended implication, as the perspectives and models align conceptually. It is reasonable to assume that the spread of a cat bond, much like returns in the stock market or insurance contracts, incorporates frictional costs, denoted by α. These friction costs capture expenses related to agency, supervision, liquidity, and other operational factors (e.g., Froot and Stein, 1998; Zanjani, 2002). Additionally, it can be assumed that beta-risk factors also contribute to the spread, although these factors are either unknown or challenging to quantify. This perspective aligns with the insights of Barberis and Huang (2008), and offers a unifying way by defining the spread as:

$$s = \alpha + \text{beta factor} + \frac{\mathbb{E}_{\mathbb{Q}}[\bar{L}_{bond}]}{S}, \tag{4}$$

and the expected loss ratio (elr) of the bond, a metric typically disclosed by the cedent, is defined as:

$$elr = \frac{\mathbb{E}_{\mathbb{P}}[\bar{L}_{bond}]}{S}.$$

This suggests that the proposed model integrates classical factor approaches with the new insights presented in this paper.

Alternative models

To compare the proposed approach the two best-performing models from Galeotti et al. (2013) are considered. Galeotti et al. (2013) evaluates several models for their applicability to cat bonds. One of the recommended models is a simple linear regression that uses only the expected loss as the key explanatory variable. This model is defined as:

$$s = \alpha + \beta_{lin} \cdot \mathbb{E}_{\mathbb{P}}[\bar{L}_{bond}] + \epsilon \tag{5}$$

The second model, which aligns more closely with the proposed approach, is based on a measure transformation by Wang, interpreted as a probability distortion (Wang, 2004). Here, Q_k denotes the Student's t-distribution with k degrees of freedom, and ϕ is the cdf of the standard normal distribution. $PFL = \mathbb{P}(\bar{L} \geq T)$ represents the probability of the first loss, while $PLL = \mathbb{P}(\bar{L} \geq K)$ denotes the probability of the last loss. The regression is expressed as:

$$s = \frac{1}{2} \cdot \Big[Q_k(\phi^{-1}(PFL) + \lambda_W) + Q_k(\phi^{-1}(PLL) + \lambda_W) \Big] + \epsilon_{\lambda_W}, \tag{6}$$

with the regression parameters $k \in \mathbb{N}$ and λ_W. Although this model is conceptually similar to the approach presented here, a key difference lies in the treatment of the variables PFL and PLL. In the alternative model, these variables are included statically in the regression. By contrast, the proposed model incorporates the full dynamics of the underlying jump processes.

3 Empirical evidence of a tail overweighting

Institutional background

The issuance of cat bonds began in the late 20th century, marking a significant development in the securitization of risk. According to data from the Artemis platform (Artemis, 2025a), the cat bond market has expanded considerably over time, both in size and complexity, encompassing an increasingly diverse array of risk perils (see Table 1). Between December 2023 and November 2024, 110 cat bonds were issued, totaling 19,697 million USD in volume. This represents a steady increase compared to the previous two periods, where 83 bonds (13,882 million USD) and 80 bonds (13,383 million USD) were issued, respectively. A defining feature of the cat bond market is its heterogeneity. In the most recent period, 61 distinct risk perils were issued, including 28 new risk perils not observed in the previous two periods. However, only 15 risk perils were issued more than twice, and just seven were issued more than three times. This means that over three-quarters of the risk perils were associated with a single bond, underscoring the diversity within the market. Such heterogeneity presents challenges for risk-based calibration, causing prior research to group thematic similar risks.

For example, U.S. hurricane risks are often aggregated at the national level rather than by individual states.

	Dec. 21 - Nov. 22	Dec. 22 - Nov. 23	Dec. 23 - Nov. 24
Bonds issued	80	83	110
Volume in Mio. USD	13,383	13,882	19,697
Risk perils	44	53	61
Unique risk perils	24	19	28
Risk perils issued at least 2 times	12	14	15
Risk perils issued at least 3 times	6	7	7

Table 1: Overview of the cat bond market in the last three periods. Source: Artemis (2025a).

This table provides an overview of the cat bond market over the last three periods, based on data from Artemis (2025a). It shows, for each period, the number of bonds issued, the total volume in million USD, and the various risk perils. Additionally, it highlights the number of unique risk perils, indicating how many risk perils were not issued in any of the other periods. The last two rows present the number of bonds issued with the same risk perils at least twice and three times, respectively.

Sharing the same risk peril does not necessarily imply identical underlying risks, as risk perils often bundle different kinds of risks. This means that even bonds categorized under the same peril may not cover homogeneous risks or share the same cdf. For instance, within the risk perils issued at least three times, some types, such as Florida and North Carolina named-storms, are considered homogeneous due to their localized geographical scope and predictable nature. In contrast, types like "unknown property catastrophe risks" and "mortgage insurance risks" exhibit a higher degree of heterogeneity (see Table 2). Cyber risk, a more recent addition to the market, remains homogeneous, as current bonds do not differentiate between distinct cyber incidents.

The transfer of cyber risk is a noteworthy development in the cat bond market (AON, 2024). Since January 2023, when Beazley issued the first bond, nine additional bonds with the underlying risk peril "cyber" have been issued by four more cedents. This group currently exhibits limited heterogeneity, with most bonds being indemnity-triggered, having primarily been issued since end 2023, and categorized broadly under cyber risks. However, Hannover Re's smaller 13.75 million Cumulus Re bond, focusing specifically on cloud-outages, represents an exception. This relative homogeneity offers a unique opportunity to analyze investor risk attitudes within a single period (see Table 3). Based on industry feedback, the cyber market is expected to grow significantly in the near future, potentially surpassing soon the billion-dollar mark. Additionally, it is anticipated that the risk heterogeneity of the bonds will increase, leading to the issuance of more specialized bonds, such as Hannover Re's cloud-outage bond. This shift would align cyber with broader trends in the cat bond market, complicating analyses that rely on a joint cdf and highlighting the timing of the analysis

Risk peril	# Bonds	Volume in Mio. USD	Classification
Unknown property catastrophe risks	20	454.47	Heterogeneous
U.S. property catastrophe risks	8	94	Heterogeneous
Cyber risks	6	785	Homogeneous
Florida named-storm	5	2000	Homogeneous
Mortgage insurance risks	3	774.6	Heterogeneous
North Carolina named-storm	3	995	Homogeneous
U.S., Canada named-storms and earthquakes	3	330	Heterogeneous

Table 2: Overview of the risk perils issued more than three times the last three period. Source: Artemis (2025a).

This table presents information on risk perils issued more than three times between December 2023 and November 2024, based on data from Artemis (2025a). It includes the number of bonds issued, the volume in million USD, and a risk classification. A homogeneous classification indicates that the bonds within the risk peril share the same cdf.

here.

Bond	Date of issue	Modeling agency	Cedent	Attachment point	Expected loss ratio	Spread
1	Nov 2023	CyberCube	AXIS Capital	2.46%	1.97%	9.75%
2	Dec 2023	CyberCube	Swiss Re	2.228%	1.721%	12%
3	Dec 2023	CyberCube	Chubb	2.142%	1.387%	9.25%
4	Dec 2023	RMS	Beazley	1.71%	1.26%	13-13.25%
5	Sep 2024	RMS	Beazley	1.2%	0.93%	10.5%

Table 3: Public information of the target bonds.

This table presents public information on bonds with the underlying risk peril "cyber" sourced from Artemis (2025a). For five bonds, it includes the date of issuance, the corresponding modeling agency, and the cedent. Additionally, bond-specific details such as the attachment quantile, the expected loss as a percentage of the notional amount, and the spread as a percentage of the notional amount are also provided.

Data

Previous research has predominantly combined bonds with different risks, aiming to increase data availability. In contrast, this paper adopts the opposite approach, emphasizing the significant role of the individual cdf in pricing. While earlier studies have attempted to account

for distribution characteristics through metrics such as expected loss or the probability of first or last loss, these methods tend to be static in nature. Effective calibration requires the cdf to be accurately aligned with the underlying distribution of the bonds. However, this presents a considerable challenge due to the scarcity or absence of long-term data. For instance, no cyber event has yet triggered a bond, making it nearly impossible to calibrate distribution parameters and jump processes without relying on industry expertise. This lack of historical data also increases the risk of confounding factors, such as the peso problem, complicating the calibration further. This paper aims to address the mentioned issues by using unique cyber data. The dataset is obtained from a leading risk modeling agent. Based on a portfolio of 10,000 policies from a global insurance broker, which represents the worldwide cyber risk market, the model agency calibrated 10,000 losses per policy, ranging from daily incidents to extreme events. The resulting 100 million data points also contain information about correlations and cross-dependencies. The individual loss simulations can accordingly be compressed into common loss distribution, which consists of 10,000 annual possible cyber market loss scenarios. Since it is known which loss count towards daily and which towards extreme, it is possible to precisely calibrate the respective components of the loss distribution from equation 2.[2] Since the modeling agency was also involved in several bond transactions with the risk peril "cyber", it can be assumed that this is the best estimate of the underlying loss distribution achieved to date.

From the cat bond market, the analysis includes five market-traded bonds with the underlying risk peril "cyber" from four different cedents, as shown in Table 3.[3] As a result, the analysis encompasses all cyber risk bonds issued during the given period, covering over 83% of all bonds ever issued and including every cedent at least once. This analysis encompasses around 4% of the total volume of new cat bond market transactions globally and nearly 6% of the transactions conducted during this period.

Calibration

The bond characteristics are initially calibrated according to the underlying cdf. Starting from the attachment quantile T, the end quantile K is determined to match the expected loss ratio, defined as:

$$elr = \frac{\mathbb{E}_{\mathbb{P}}[\bar{L}_{bond}(T,K)]}{K-T}.$$

Once the quantiles for the bond are determined, the value of η is searched which satisfies:.

$$s = \frac{\mathbb{E}_{\mathbb{Q}}[\bar{L}_{bond}(T,K)]}{K-T}.$$

Table 4 presents the calibration results, including the attachment and end quantiles, the multiple, and the estimated η for each bond.

[2]However, the specific calibration results cannot be provided here due to the confidentiality of the data. Further details about the dataset are available in Cremer et al. (2025).

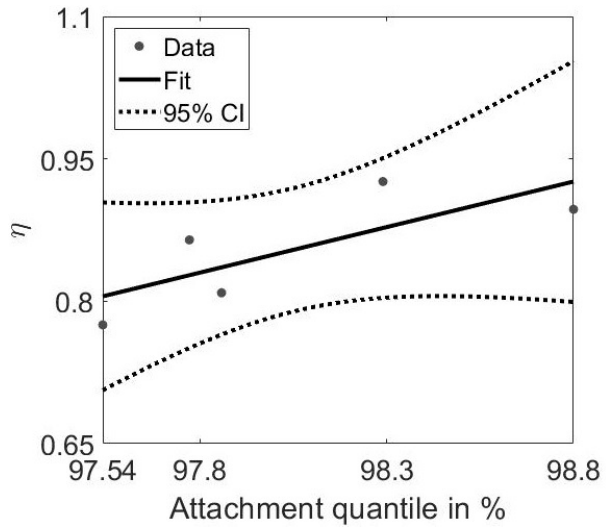
[3]Strictly speaking, the calibration relies on six bonds. Bond 4 from Table 3 represents two separate issuances that emerged from a quota share process. These bonds were issued in December 2023 and May 2024, with approximately 45% of the loss allocated to the first bond and the remaining 55% to the second.

Bond	Attachment quantile	End quantile	Multiple	η
1	97.54%	99.583%	4.9492	0.7411
2	97.772%	99.617%	6.9727	0.8454
3	97.858%	99.697%	6.6691	0.8060
4	98.29%	99.699%	10.3175 to 10.5159	0.9372 to 0.9439
5	98.8%	99.77%	11.29	0.9370

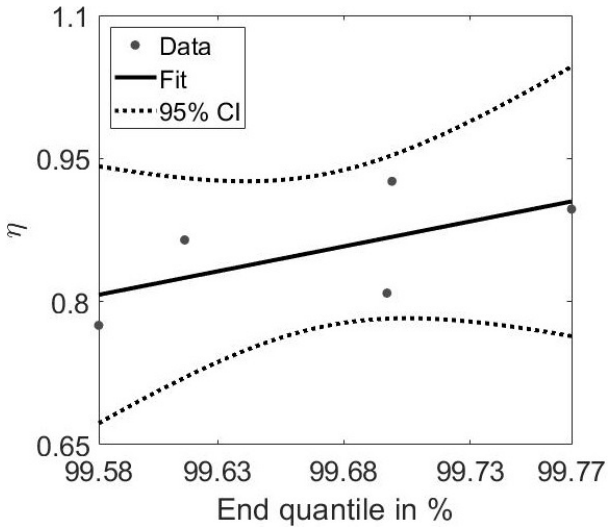
Table 4: Calibration output

This table provides calibration information for the five bonds. It includes the attachment quantile and end quantile as percentages, along with the computed bond multiple. The final column presents the estimated η, which represents the transformation parameter of the $\mathbb{Q}$ measure.

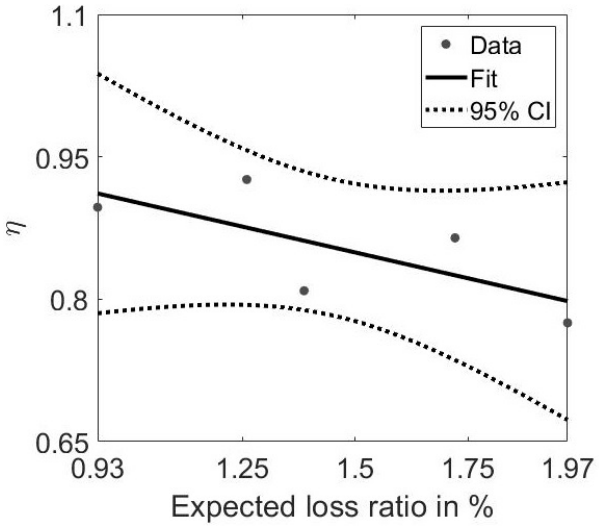
Next, the relationships between bond characteristics and η are examined. Following Galeotti et al. (2013), the focus is on three characteristics: expected loss, probability of first loss, and probability of last loss. A simple linear relationship is assumed for this analysis. The pattern in Figure 2 does not indicate that this is an insufficient assumption. Figure (a) presents the best fit with an R^2 of 0.59, illustrating the relationship between the attachment quantile, representing the probability of first loss, and transformation parameter. The upwards slope indicates that the higher the attachment quantile, meaning the lower the probability of triggering the bond, the more the jump process is overemphasized. The relationship with the end quantile, representing the probability of last loss, shown in Figure (b), is weaker, with an R^2 of 0.39. However, a similar trend is observed: the lower the probability of occurrence, the more the jump is overweighted. Figure (c) shows the relationship with the expected loss ratio. Here, a higher expected loss corresponds to lower overweighting of the jump process. This relationship achieves an R^2 of 0.51. Low-frequency, high-severity events typically have low expected loss ratios since higher expected loss ratios tend to cover losses earlier in the cdf, capturing proportionately less tail risk. Therefore, the tendency from all three calibrations is clearly recognizable: the more risk from the jump process is proportionally in the bond, the more the jump is overestimated.



(a) Attachment quantile with $R^2 = 0.59$.



(b) End quantile with $R^2 = 0.39$.



(c) Expected loss ratio with $R^2 = 0.51$.

Figure 2: Relationships with η.

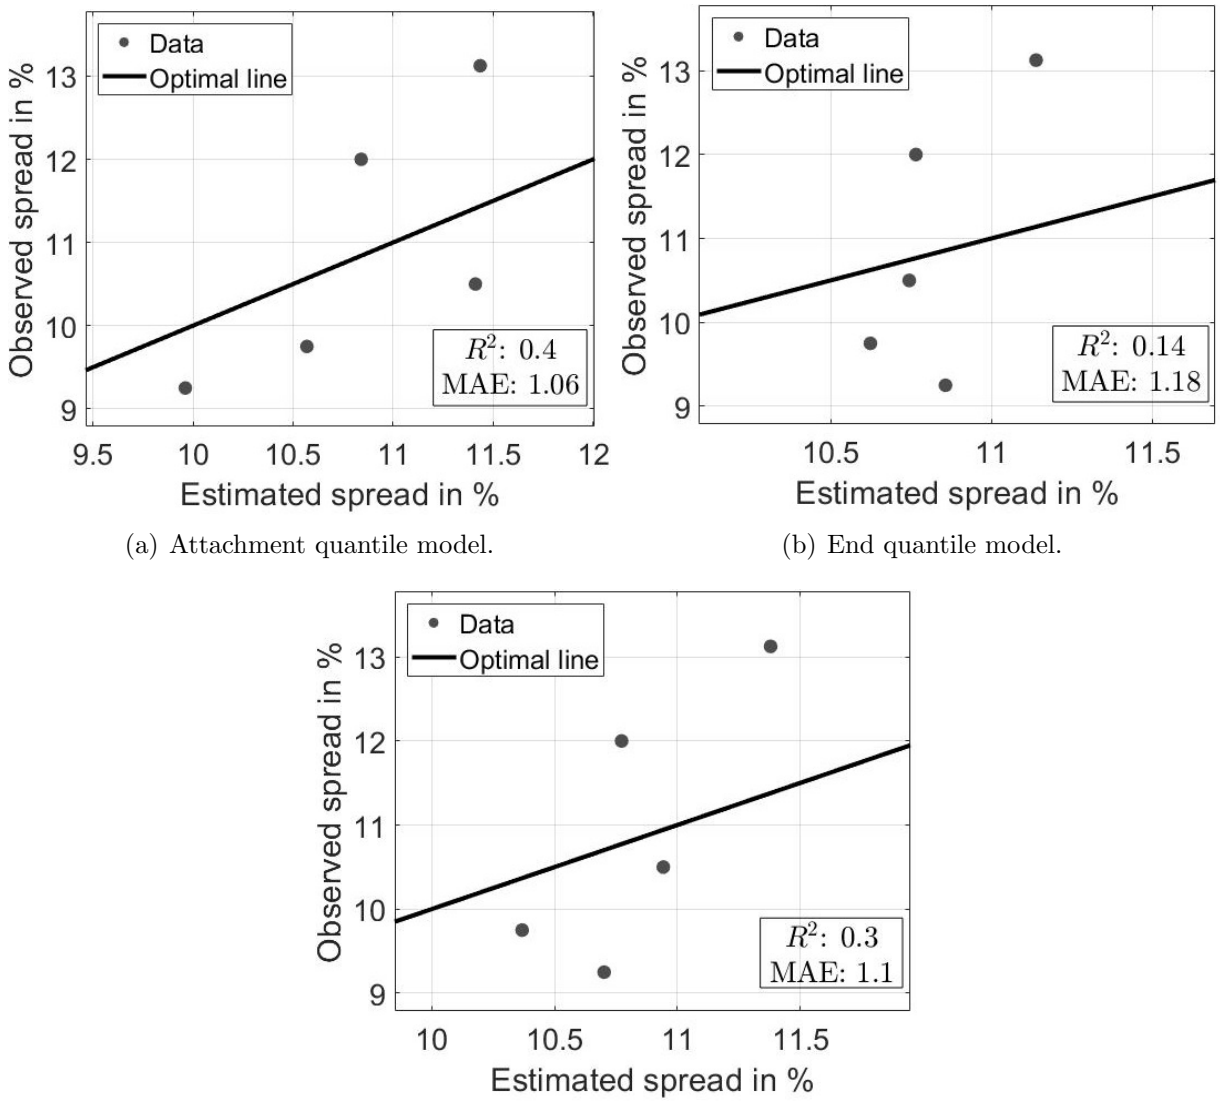
This figure presents a linear regression fit of (a) the attachment quantile, (b) the end quantile and (c) the expected loss ratio with the transformation parameter η. The data are in points, the solid line is the regression fit, and the dashed line is the 95% confidence internal (CI).

Although the attachment quantile model provides the best fit, all three models will be considered in the next step. Based on the respective linear model, spreads are calculated for the input bonds. Figure 3 illustrates the goodness of fit: (a) for the attachment quantile model, (b) for the end quantile model, and (c) for the expected loss ratio model.

It is evident that the end quantile model is not suitable for estimating η and the spread. This model previously had the lowest R^2 and now performs worst in terms of both R^2 and mean absolute error (MAE). The pattern of points, which are more vertically than horizontally aligned around the optimal line, indicates a poor model fit. The model based

on the expected loss ratio performs slightly better. While its R^2 and MAE values improve, the data points still do not align well horizontally around the optimal line.

In contrast, the attachment quantile model shows the best performance in estimating η and in replicating the spreads. With an explanatory power of 40% for the spread, it is the strongest model. More importantly, the estimates are well distributed around the optimal line, further confirming its superior fit.



(a) Attachment quantile model.

(b) End quantile model.

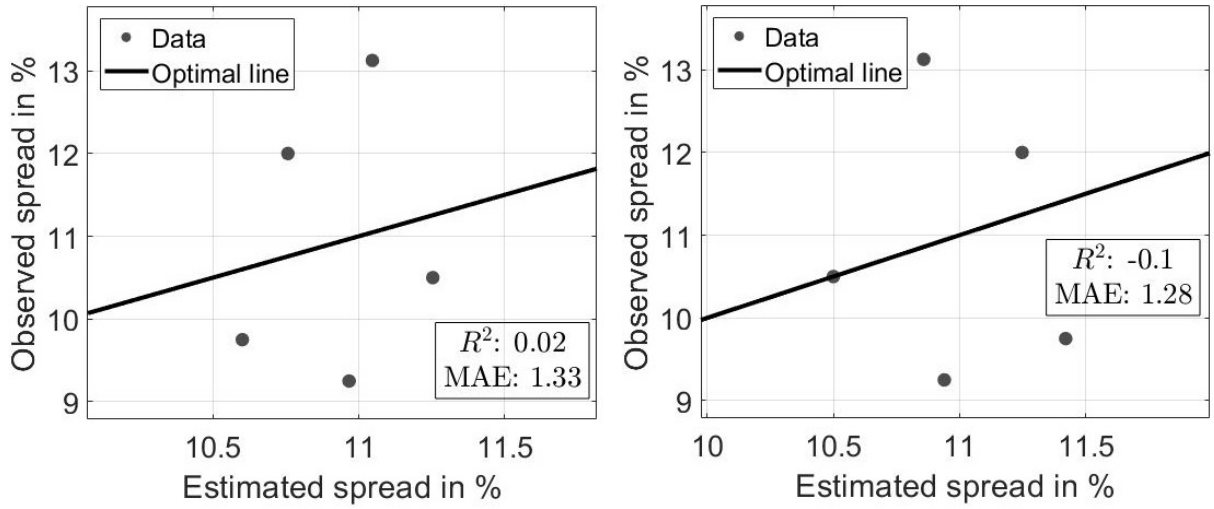
(c) Expected loss ratio model.

Figure 3: Model predictions.

This figure presents the model predictions of (a) the attachment quantile model, (b) the end quantile model and (c) the expected loss ratio model. The data are in points, the solid line is the optimal line defining a perfect fit. Furthermore, the R^2 and the mean absolute error (MAE) are given.

Alternative models

To provide context for the 40% explanatory power, a linear model was introduced in equation 5, and a model based on the Wang transformation was presented in equation 6. According to Galeotti et al. (2013), these two models performed best in a cross-sectional setting when compared to various other approaches. However, when applied in this setting, both models exhibit virtually no explanatory power (Figure 4). This result is not unexpected, as the analysis here focuses on a homogeneous rather than a heterogeneous set of bonds. At first glance, the lack of explanatory power in the Wang model – despite its ability to extend beyond standard distributional assumptions – might seem surprising. However, this outcome actually reinforces the perspective taken here. It suggests that, in a cross-sectional context, the Wang model effectively filters out variations in η across different risk perils (and time periods). As a result, it has no explanatory power in the homogeneous environment. A similar reasoning applies to the insignificance of the linear model. Although these models do not provide insights in this particular setting, the cross-sectional analysis of the cat bond market remains highly relevant. This will be further demonstrated in the following.



(a) Linear model from equation 5.

(b) Wang model from equation 6.

Figure 4: Alternative model predictions.

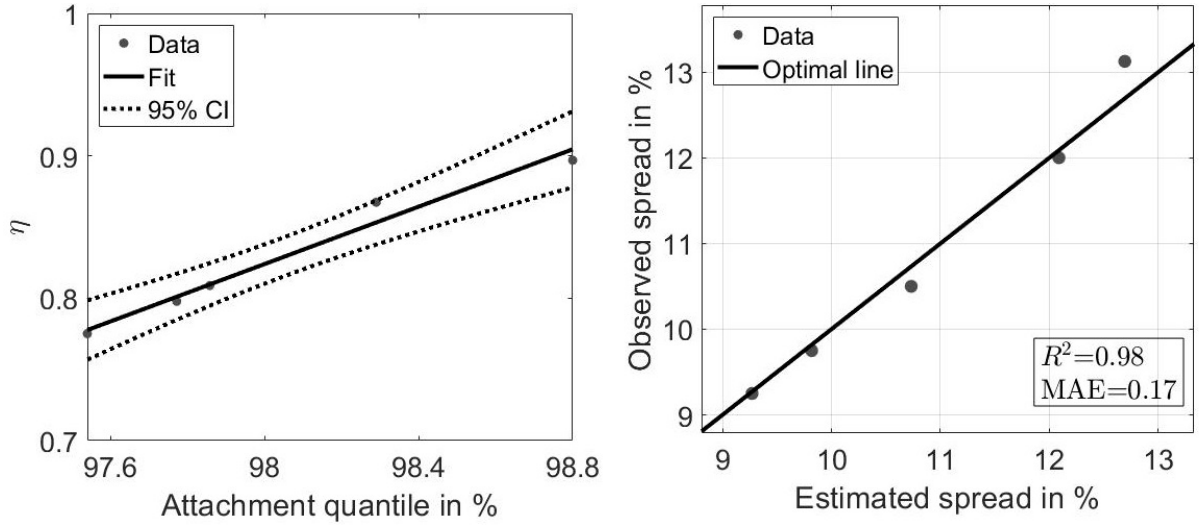
This figure presents alternative model predictions of (a) the linear model (equation 5) and (b) the Wang model (equation 6). The data are in points, the solid line is the optimal line defining a perfect fit. Furthermore, the R^2 and the mean absolute error (MAE) are given.

Beta factors & tail overweighting

The fit shown in Figure 3 (part a) indicates that two spreads are underestimated, while three are overestimated. Since the estimated values in these two groups are almost equally distant from the optimal line, it can be assumed that bond-specific risk premiums exist – particularly for the two underestimated bonds – which are not yet be accounted for. As shown in equation 4, such premiums represent components of the spread that can be attributed to

beta factors. The idea that certain bonds exhibit a higher spread than comparable ones is not unreasonable. For example, Braun's empirical research demonstrates that cedents – such as a cat bond issued by Swiss Re – can differ significantly from others (Braun, 2016). These differences are primarily cross-sectional in nature. Suppose bonds 2 and 4 from the Table 3 have a beta factor of 2%. The same estimation procedure is then applied, as shown in Figure 5. First, η is estimated using a linear model with the attachment quantile as an independent variable (part a). This model is then used to estimate the spread (part b). The results indicate that the tail overweighting combined with the beta factor can explain the spread almost perfectly.[4]

Summarized, the findings from both this analysis and the literature suggest that the models used for estimating spread – whether cross-sectional or cdf-specific – offer a certain degree of explanatory power. However, a significant portion of the variation remains unexplained. The calibration here demonstrates that integrating both modeling approaches can yield an almost perfect result.



(a) Relationship with η; $R^2 = 0.97$.

(b) Model prediction.

Figure 5: Calibration with beta factors.

This figure presents calibration with beta factors. (a) is the relationship of the attachment quantile to η, (b) is the spread estimation with the R^2 and the mean absolute error (MAE). In part (a), the data are in points, the solid line is the regression fit, and the dashed line is the 95% confidence internal (CI).

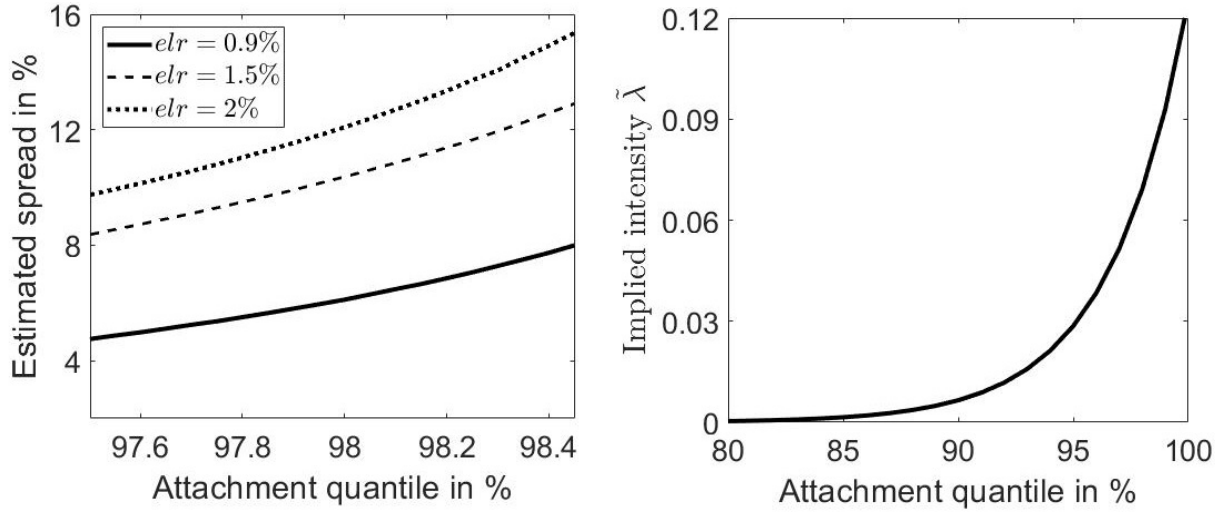
Economic interpretation

The results indicate that the higher the attachment quantile – meaning the lower the probability that the bond will be triggered – the higher the relative risk premium investors expect (tail overweighting). This pattern is further illustrated in Figure 6, part (a), which presents

[4]The scenarios are also tested with alternative beta factor values as well as with different frictions ($\alpha > 0$). The conclusion that a combined approach leads to high explanatory power remains consistent across these variations, see Appendix A.2.

spread calculations for synthetic bonds. Three expected loss ratios are first fixed, and bonds with varying attachment quantiles are calibrated using the industry data. In a world where $\mathbb{Q}$ and $\mathbb{P}$ are identical, the spread would match the expected loss ratio. However, the results show that the estimated spread steadily increases with higher attachment quantiles.

This behavior cannot be explained solely by typical investor risk aversion, as prior studies have shown. Barberis and Huang (2008) demonstrate that a security with positively skewed returns tends to be overpriced relative to expectations in a traditional beta-risk-based framework. The pricing measure $\mathbb{Q}$ in this paper's model captures precisely this overpricing. Moreover, the relationship between $\mathbb{Q}$ and tail risk indicates that relative overpricing increases as tail risk grows. Notably, this tendency aligns with observations in the insurance market. Sydnor (2010) show that under probability weighting, households overemphasize tail events. As a result, their heightened focus on these rare but highly undesirable outcomes leads them to pay a higher premium for insurance policies. This behavior has a well-known parallel in the options market: the volatility smile (e.g., Derman and Kani, 1994). In standard option pricing models, deep in-the-money or out-of-the-money options exhibit strike price deviations that are not fully captured by traditional models (analogous here to beta factors). To better align with market prices, these models incorporate implied volatility, which, in this context, corresponds to the pricing measure $\mathbb{Q}$. Similarly, the overweighting observed here follows the same principle by defining an implied intensity (Equation 3) over the transformation parameter η. However, due to the one-sided skewness of the catastrophe bond market distribution, only a half-smile emerges, as shown in Figure 6, part (b).



(a) Estimated spread.

(b) Estimated implied intensity.

Figure 6: Half-smile in cat bond pricing.

This figure illustrates the half-smile pattern in catastrophe bond spreads, based on the η estimate from Figure 5. Part (a) presents three expected loss ratios, with gradually increasing attachment quantiles, extending the bond's exposure to jump risk. Part (b) estimates the implied intensity using equation 3 on the same basis. As a result, both the spread and implied intensity increase, forming a half-smile pattern.

Robustness, limitations & future research

Three main limitations can be identified in this approach. The first concerns the number of bonds inside the calibration. While this study takes a fundamentally different approach from cross-sectional analyses by focusing on a homogeneous set of bonds based on the cdf, the sample size remains limited to five bonds. Although this is not ideal, cyber risk remains a suitable choice for this analysis. As shown in Table 1, only Florida named-storm bonds offer a comparable number of homogeneous instruments, and no other risk peril provides a larger calibration sample. Furthermore, Table 5 examines whether different permutations fall within the 95% confidence interval of the five-bond calibration. The model was evaluated both with and without the beta factor. The findings indicate that a calibration based on just three bonds are in 95% of the cases inside the confidence range of the five-bond calibration, suggesting a high degree of stability in tail overweighting.

	Two bonds	Three bonds	Four bonds
Without beta factor:			
Permutations inside 95% CI	6	10	5
Permutations outside 95% CI	4	0	0
With beta factor:			
Permutations inside 95% CI	7	9	5
Permutations outside 95% CI	3	1	0

Table 5: Robustness bond amount.

This table provides an overview of the calibration process's stability. The calibration is performed using two, three, or four bonds, considering all possible permutations. The upper section verifies whether the calibrations fall within the 95$ confidence interval (CI) of the five-bond calibration without a beta factor from Figure 3, part (a). The lower section verifies whether the calibrations fall within the 95$ confidence interval (CI) of the five-bond calibration with the beta factor from Figure 5, part (a).

The second limitation is that the analysis is restricted to cyber risk. Although Florida named-storm bonds offer a similarly homogeneous sample, there is insufficient data for a robust distribution calibration. Additional distributions were not approved by the model agency, and independent estimation is practically infeasible due to data constraints. Even if a Florida named-storm loss distribution is available, it would lack insights into the covered losses. Understanding the bond's loss dynamics requires a representative sample of policies for the region, similar to the calibration provided by the global cyber broker. This highlights the uniqueness of the present analysis, which offers a rare insight on these dynamics. Nevertheless, Appendix A.3 examines Florida named-storms and applies tail overweighting. The results indicate a similar overweighting pattern, with a combination of cdf-specific and beta factor approaches replicating spreads almost perfectly. Despite the limited data availability,

this pattern cannot be disregarded as coincidence.

The third limitation relates to the choice of distribution. While the compounded Poisson process is widely applied in tail risk modeling, alternative approaches should also be mentioned. One example is a mixture model combining an empirical distribution with a generalized Pareto tail (e.g., McNeil et al., 2015), which is frequently used in the context of natural disasters (e.g., Braun et al., 2023). However, it is crucial to ensure that transformations such as the Esscher transformation (or similar methods) exist and remain well-defined. In the case of fat-tailed distributions with infinite moments, this condition is often not met, posing a significant challenge.

Labeled as limitations, these points also present valuable opportunities for future research. One important direction is to examine how tail overweighting behaves across different risk perils and time periods. For example, the Florida and South Carolina named-storm bond issued by Slide Insurance in April 2024 shares many characteristics with the Beazley bond issued in September 2024, yet its spread is 1.5% lower. Understanding if such differences arrives from beta factors or risk perceptions could provide deeper insights into market pricing mechanisms. Additionally, this research highlights the need to incorporate tail overweighting into future cross-sectional models or to develop methods that better integrate these approaches. Further work could explore techniques to align these methodologies more effectively. Finally, drawing parallels to the volatility smile in options markets opens another direction for investigation. Future studies could analyze whether similar (half-)smiles appear in other financial markets and assess their implications for pricing.

4 Concluding remarks

This paper introduces a new cat bond pricing perspective by incorporating tail overweighting in a jump-risk framework. Traditional factor models, by design, fail to fully capture cat bond spreads because they cannot account for the idiosyncratic nature of extreme event risks. By leveraging a calibrated measure transformation, this study demonstrates that investors systematically overweight tail components, particularly for bonds with low attachment probabilities. The empirical results for the risk peril "cyber" indicate that approximately 40% of the spread can be attributed to tail overweighting alone. When combined with cross-sectional beta factors, this approach enables near-perfect replication of observed spreads, suggesting that cat bond pricing is best understood through an integrated framework rather than separate models for systematic and cdf-driven risks.

The findings contribute to broader asset pricing research by highlighting the significance of probability distortions in markets exposed to extreme downside risk. Similar distortions have been observed in options markets and household insurance, underscoring the need for pricing models that account for investor tail-risk aversion. The observed probability distortions align with behavioral finance theories, such as probability weighting, and suggest that cat bond investors price tail risk similarly to option traders who incorporate implied volatility adjustments. But certain limitations remain. The analysis focuses on a homogeneous subset of cat bonds – the risk peril "cyber" – which provides a controlled setting for calibration but limits the generalizability across broader cat bond markets. Future research

could extend this approach to other risks, such as natural disasters, where loss distributions exhibit more heterogeneity. Additionally, refining the integration of tail overweighting into cross-sectional asset pricing models could provide further insights into risk premiums across financial markets.

Overall, this study enhances the understanding of cat bond pricing by bridging actuarial probability distortions with financial market pricing mechanisms. By demonstrating the explanatory power of tail overweighting in a structured setting, the findings offer a valuable foundation for improving risk assessment, capital allocation, and market efficiency in cat bonds and beyond.

References

Andersen, T. G., Benzoni, L., and Lund, J. (2002). An empirical investigation of continuous-time equity return models. *Journal of Finance*, 57(3):1239–1284.

AON (2024). ILS Annual Report. https://www.aon.com/getmedia/154b74d4-b861-45a5-a14c-bc258c88d19f/20240830-ils-annual-report-2024.pdf&ved=2ahUKEwjjv4jY58qLAxWgh_0HHTozAQQQFnoECAoQAQ&usg=AOvVaw0FQOgFOwj2mSO4ccuXzc_K. [Online; accessed 09-February-2025.

Artemis (2025a). Catastrophe Bond & Insurance-Linked Securities Deal Directory. https://www.artemis.bm/deal-directory. [Online; accessed 09-February-2025.

Artemis (2025b). Catastrophe bonds & ILS outstanding by trigger type. https://www.artemis.bm/dashboard/cat-bonds-ils-by-trigger/. [Online; accessed 09-February-2025].

Artemis (2025c). Catastrophe bonds and ILS average multiple by year. https://www.artemis.bm/dashboard/cat_bonds_ils_average_multiple/. [Online; accessed 09-February-2025].

Artemis (2025d). What is a catastrophe bond (or cat bond)? https://www.artemis.bm/library/what-is-a-catastrophe-bond/. [Online; accessed 09-February-2025].

Bai, J., Bali, T. G., and Wen, Q. (2019). Common risk factors in the cross-section of corporate bond returns. *Journal of Financial Economics*, 131(3):619–642.

Bakshi, G. and Madan, D. (2002). Average rate claims with emphasis on catastrophe loss options. *Journal of Financial and Quantitative Analysis*, 37(1):93–115.

Barberis, N. and Huang, M. (2008). Stocks as lotteries: The implications of probability weighting for security prices. *American Economic Review*, 98(5):2066–2100.

Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *The Review of Financial Studies*, 9(1):69–107.

Beer, S. and Braun, A. (2022). Market-consistent valuation of natural catastrophe risk. *Journal of Banking & Finance*, 134:106350.

Ben-Ammar, S., Eling, M., and Milidonis, A. (2018). The cross-section of expected stock returns in the property/liability insurance industry. *Journal of Banking & Finance*, 96:292–321.

Bollerslev, T. and Todorov, V. (2011). Tails, fears, and risk premia. *Journal of Finance*, 66(6):2165–2211.

Bowers, N. L., Gerber, H. U., Hickman, J. C., Jones, D. A., and Nesbitt, C. (1986). *Actuarial Mathematics*. Itasca, Ill.: Society of Actuaries.

Braun, A. (2016). Pricing in the primary market for cat bonds: New empirical evidence. *Journal of Risk and Insurance*, 83(4):811–847.

Braun, A., Ammar, S. B., and Eling, M. (2019). Asset pricing and extreme event risk: Common factors in ILS fund returns. *Journal of Banking & Finance*, 102:59–78.

Braun, A., Eling, M., and Freyschmidt, M. (2023). Optimal government intervention for extreme risks. *Available at SSRN*.

Cremer, F., Freyschmidt, M., Kasper, D., and Ning, D. (2025). On the measurement of cyber risk: A comparative analysis of public and industrial data. *Working Paper*.

Cummins, J. D. and Phillips, R. D. (2005). Estimating the cost of equity capital for property-

liability insurers. *Journal of Risk and Insurance*, 72(3):441–478.

Cummins, J. D. and Weiss, M. A. (2009). Convergence of insurance and financial markets: Hybrid and securitized risk-transfer solutions. *Journal of Risk and Insurance*, 76(3):493–545.

Derman, E. and Kani, I. (1994). Riding on a smile. *Risk*, 7(2):32–39.

Dickerson, A., Mueller, P., and Robotti, C. (2023). Priced risk in corporate bonds. *Journal of Financial Economics*, 150(2):103707.

Doherty, N. A. and Garven, J. R. (1986). Price regulation in property-liability insurance: A contingent-claims approach. *Journal of Finance*, 41(5):1031–1050.

Eraker, B. (2004). Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *Journal of Finance*, 59(3):1367–1403.

Esscher, F. (1932). On the probability function in the collective theory of risk. *Skandinavisk Aktuarietidskrift*, 15(3):175–195.

Froot, K. A. and Stein, J. C. (1998). Risk management, capital budgeting, and capital structure policy for financial institutions: An integrated approach. *Journal of Financial Economics*, 47(1):55–82.

Galeotti, M., Gürtler, M., and Winkelvos, C. (2013). Accuracy of premium calculation models for cat bonds – An empirical analysis. *Journal of Risk and Insurance*, 80(2):401–421.

Gao, G. P., Lu, X., and Song, Z. (2019). Tail risk concerns everywhere. *Management Science*, 65(7):3111–3130.

Gerber, H. U. and Shiu, E. S. (1994). Option pricing by Esscher transforms. *Transactions of Society of Actuaries*, 46:99–191.

Grinsted, A., Ditlevsen, P., and Christensen, J. H. (2019). Normalized US hurricane damage estimates using area of total destruction, 1900-2018. *Proceedings of the National Academy of Sciences*, 116(48):23942–23946.

Gürtler, M., Hibbeln, M., and Winkelvos, C. (2016). The impact of the financial crisis and natural catastrophes on cat bonds. *Journal of Risk and Insurance*, 83(3):579–612.

Harrison, J. M. and Pliska, S. R. (1981). Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and their Applications*, 11(3):215–260.

Herrmann, M. and Hibbeln, M. (2021). Seasonality in catastrophe bonds and market-implied catastrophe arrival frequencies. *Journal of Risk and Insurance*, 88(3):785–818.

Kai-Ineman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):363–391.

Kou, S. G. (2002). A jump-diffusion model for option pricing. *Management Science*, 48(8):1086–1101.

Lee, J.-P. and Yu, M.-T. (2002). Pricing default-risky cat bonds with moral hazard and basis risk. *Journal of Risk and Insurance*, 69(1):25–44.

Ma, Z.-G. and Ma, C.-Q. (2013). Pricing catastrophe risk bonds: A mixed approximation method. *Insurance: Mathematics and Economics*, 52(2):243–254.

McNeil, A. J., Frey, R., and Embrechts, P. (2015). *Quantitative risk management: Concepts, techniques and tools-Revised edition*. Princeton university press.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1-2):125–144.

Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated

time-series study. *Journal of Financial Economics*, 63(1):3–50.

Sydnor, J. (2010). (Over)insuring modest risks. *American Economic Journal: Applied Economics*, 2(4):177–199.

Tomunen, T. (2022). Failure to share natural disaster risk. *Available at SSRN 3525731*.

Wang, S. S. (2004). Cat bond pricing using probability transforms. *Geneva Papers: Etudes et Dossiers*, 278:19–29.

Zanjani, G. (2002). Pricing and capital allocation in catastrophe insurance. *Journal of Financial Economics*, 65(2):283–305.

A Appendix

A.1 Measure Changes using the Esscher Transformation

Following Gerber and Shiu (1994), let be:

$$L_t = L_0 \exp(X_t),$$

where X_t with $t \geq 0$ is a stochastic process characterized by stationary and independent increments, and $X_0 = 0$. Furthermore, let:

$$F_{X_t}(x) = \mathbb{P}(X_t \leq x)$$

be the cdf, and:

$$M_{\mathbb{P},X_t}(u) = \mathbb{E}[\exp(uX_t)]$$

represent the moment-generating function of X_t under the measure $\mathbb{P}$. Thus:

$$M_{\mathbb{P},X_t}(u) = \int_{-\infty}^{\infty} \exp(ux) f(x,t) dx,$$

where $f(x,t)$ is the continuous density of X_t. Building upon the transformation proposed by Esscher (1932), a transformed density for X_t is (Gerber and Shiu, 1994):

$$\begin{aligned} f(x,t,h) &= \frac{\exp(hx) f(x,t)}{\int_{-\infty}^{\infty} \exp(hy) f(y,t) dy} \\ &= \frac{\exp(hx) f(x,t)}{M_{\mathbb{P},X_t}(h)} \end{aligned}$$

where h is the transformation parameter. The corresponding moment-generating function is given by (Gerber and Shiu, 1994):

$$\begin{aligned} M_{\mathbb{Q},X_t}(u) &= \int_{-\infty}^{\infty} \exp(ux) f(x,t,h) dx \\ &= \frac{M_{\mathbb{P},X_t}(u+h)}{M_{\mathbb{P},X_t}(u)}. \end{aligned}$$

Subsequently, the Esscher transformation is derived for the three significant processes in this study, aligning with Gerber and Shiu (1994). In the provided examples, the time component is disregarded.

Normal Distribution: Assuming $X = Y$, where Y is a normally distributed random variable with a mean of μ and a variance of σ^2. The moment-generating function under $\mathbb{P}$ is expressed as:

$$M_{\mathbb{P},X}(u) = \exp(u\mu + \frac{1}{2}\sigma^2 u^2).$$

Through the Esscher transformation, the resulting expression for the moment-generating function under the new measure $\mathbb{Q}$ is (Gerber and Shiu, 1994):

$$M_{\mathbb{Q},X}(u) = \exp\Big(u(\mu + h\sigma^2) + \frac{1}{2}\sigma^2 u^2\Big).$$

Consequently, the new mean under $\mathbb{Q}$ can be defined as $\tilde{\mu} = \mu + h\sigma^2$. The transformed normal distribution under $\mathbb{Q}$ remains a normal distribution with mean $\tilde{\mu}$ variance σ^2.

Proof.

$$\begin{aligned}
\frac{M_{\mathbb{P},X}(u+h)}{M_{\mathbb{P},X}(u)} &= \frac{\exp\left((u+h)\mu + \frac{1}{2}\sigma^2(u+h)^2\right)}{\exp(u\mu + \frac{1}{2}\sigma^2 u^2)} \\
&= \exp\Big((u+h)\mu + \frac{1}{2}\sigma^2(u+h)^2 - (u\mu + \frac{1}{2}\sigma^2 u^2)\Big) \\
&= \exp\Big((h\mu + \frac{1}{2}\sigma^2 h^2 + \sigma^2 uh\Big) \\
&= \exp\Big((h(\mu + \sigma^2 u) + \frac{1}{2}\sigma^2 h^2\Big)
\end{aligned}$$

□

Poisson Distribution: Assume $X = kN$, where N is a Poisson process with intensity λ, and k is a constant. The moment-generating function under $\mathbb{P}$ is defined as:

$$M_{\mathbb{P},X}(u) = \exp\Big(\lambda(\exp(ku) - 1)\Big)$$

Through the Esscher transformation, the resulting expression for the moment-generating function under the new measure $\mathbb{Q}$ is (Gerber and Shiu, 1994):

$$M_{\mathbb{Q},X}(u) = \exp\Big(\lambda \exp(hk)(\exp(ku) - 1)\Big)$$

Consequently, the intensity under $\mathbb{Q}$ can be defined as $\tilde{\lambda} = \lambda \exp(hk)$. The transformed Poisson process under $\mathbb{Q}$ remains a Poisson process with intensity $\tilde{\lambda}$.

Proof.

$$\begin{aligned}
\frac{M_{\mathbb{P},X}(u+h)}{M_{\mathbb{P},X}(u)} &= \frac{\exp\big(\lambda(\exp(k(u+h)) - 1)\big)}{\exp\big(\lambda(\exp(ku) - 1)\big)} \\
&= \exp\Big(\lambda(\exp(k(u+h)) - 1) - \lambda(\exp(ku) - 1)\Big) \\
&= \exp\Big(\lambda(\exp(ku)\exp(kh)) - \lambda\exp(ku)\Big) \\
&= \exp\Big(\lambda\exp(ku)(\exp(kh) - 1)\Big)
\end{aligned}$$

□

Compounded Poisson Process: Assume a compounded Poisson process $X = \sum_{i=1}^{N} Y_i$, where N is a Poisson process with intensity λ, and Y_i represents a normally distributed jump size with mean μ and variance σ^2. The moment-generating function under $\mathbb{P}$ is defined as:

$$\begin{aligned} M_{\mathbb{P},X}(u) &= \mathbb{E}[\exp(u \sum_{i=1}^{N} Y_i)] \\ &= \exp\left(\lambda(M_{\mathbb{P},Y}(u) - 1)\right) \end{aligned}$$

Through the Esscher transformation, the resulting expression for the moment-generating function under the new measure $\mathbb{Q}$ is (Gerber and Shiu, 1994):

$$M_{\mathbb{Q},X}(u) = \exp\left(\lambda M_{\mathbb{P},Y}(h)(M_{\mathbb{Q},Y}(u) - 1)\right)$$

Consequently, the intensity under $\mathbb{Q}$ can be defined as $\tilde{\lambda} = \lambda M_{\mathbb{P},Y}(h)$, and the new mean of the jump size under $\mathbb{Q}$ can be defined as $\tilde{\mu} = \mu + h\sigma^2$. The transformed compounded Poisson process under $\mathbb{Q}$ remains a compounded Poisson process with intensity $\tilde{\lambda}$, mean jump size $\tilde{\mu}$ and variance σ^2.

Proof.

$$\begin{aligned} \frac{M_{\mathbb{P},X}(u+h)}{M_{\mathbb{P},X}(u)} &= \frac{\exp\left(\lambda(M_{\mathbb{P},Y}(u+h) - 1)\right)}{\exp\left(\lambda(M_{\mathbb{P},Y}(u) - 1)\right)} \\ &= \exp\left(\lambda(M_{\mathbb{P},Y}(u+h) - 1) - \lambda(M_{\mathbb{P},Y}(u) - 1)\right) \end{aligned}$$

Given the moment-generating function of a normally distributed random variable, one obtains:

$$\begin{aligned} M_{\mathbb{P},Y}(u+h) &= \exp\left((u+h)\mu + \frac{1}{2}\sigma^2(u+h)^2\right) \\ &= \exp\left(u\mu + h\mu + \frac{1}{2}\sigma^2u^2 + \frac{1}{2}\sigma^2h^2 + \sigma^2uh\right) \\ &= \exp\left(u\mu + \frac{1}{2}\sigma^2u^2 + h(\mu + \sigma^2u) + \frac{1}{2}\sigma^2h^2\right) \\ &= M_{\mathbb{P},Y}(u)M_{\mathbb{Q},Y}(h) \end{aligned}$$

Therefore:

$$\begin{aligned} \frac{M_{\mathbb{P},X}(u+h)}{M_{\mathbb{P},X}(u)} &= \exp\left(\lambda(M_{\mathbb{P},Y}(u)M_{\mathbb{Q},Y}(h) - 1) - \lambda(M_{\mathbb{P},Y}(u) - 1)\right) \\ &= \exp\left(\lambda M_{\mathbb{P},Y}(u)M_{\mathbb{Q},Y}(h) - \lambda M_{\mathbb{P},Y}(u)\right) \\ &= \exp\left(\lambda M_{\mathbb{P},Y}(u)(M_{\mathbb{Q},Y}(h) - 1)\right) \end{aligned}$$

□

A.2 Beta factors

Section 3 demonstrates that tail overweighting, together with beta factors, can almost perfectly explain the observed spreads. For this purpose, a beta factor of 2% is assumed for bonds 2 and 4 from Table 3. This is expanded here. Figure 2, part (a) illustrates that grouping bonds 2 and 4 as well as 1, 3, and 5 is reasonable, as they lie at nearly identical levels parallel to the optimal line. The beta factors for bonds 1, 3, and 5 (and for bonds 2 and 4, respectively) are now denoted as β_{G1} (and β_{G2}, respectively). Furthermore, frictional costs α, introduced in equation 4, are assumed to be identical for all bonds. Various values are assigned to all three variables to identify patterns in explanatory power. For each set of values, the analysis from section 3 is conducted, and an R^2 is calculated between the estimated and observed spreads, as shown in Table 6. Importantly, the observed pattern consistently aligns with either Figure 3, part (a) or Figure 5, part (a), depending on whether the two groups are closer to or further from the optimal line. The linear structure remains intact.

Table 6 confirms that the near-perfect explanatory power is not solely driven by a single beta factor. As long as the identifiable structure – where bonds 2 and 4 have a higher beta factor than the other three – is maintained, nearly perfect replications of the spreads can be achieved. This result remains robust even when incorporating frictional costs or introducing beta factors for bonds 1, 3, and 5. Furthermore, the subsequent analysis for Florida named-storms exhibits the same pattern.

$\{\beta_{G1}, \beta_{G2}\}$	R^2		
	$\alpha = 0\%$	$\alpha = 1\%$	$\alpha = 2\%$
{0%, 0%}	0.40	0.37	0.34
{0%, 1%}	0.80	0.78	0.77
{0%, 2%}	0.98	0.97	0.97
{0%, 3%}	0.93	0.94	0.94
{0%, 4%}	0.66	0.68	0.70
{1%, 4%}	0.94	0.94	0.95
{2%, 4%}	0.97	0.96	0.95
{3%, 4%}	0.75	0.72	0.71
{4%, 4%}	0.27	0.23	0.18

Table 6: R^2 for different α and beta factors.

This table presents the R^2 for the estimated and observed spreads across different combinations of beta factors (β_{G1}, β_{G2}) and frictional costs (α). β_{G1} corresponds to bonds 1, 3, and 5 from Table 3, and β_{G2} to bonds 2 and 4. Frictional costs (α) are assumed to be identical across all bonds.

A.3 Florida named-storm bonds

An attempt was made to derive a cdf for named-storms in Florida using historical and normalized data. For this purpose, the data provided by Grinsted et al. (2019), which contains all named-storms in the United States from 1900 to 2018, was filtered for Florida-specific events. The data reveals that in more than 55% of the years no named storm occurred. Moreover, only one event, the Great Miami Hurricane of 1926, stands out as a significant jump. As a result, a robust calibration proves infeasible.[5]

While a jump process could, maybe, be defined using cross-sectional methods and advanced analytical tools, such an approach would be highly arbitrary and unstable due to only one observable data point. Instead, the assumption is made that tail overweighting remains sufficiently stable to allow for calibration using alternative data sources.[6] Consequently, the five Florida named-storm Class A bonds in Table 7 are calibrated to the existing cyber data. Although the assumption that the loss dynamics of Florida named-storm and cyber losses are identical is highly unlikely, this approach does not introduce greater arbitrariness than the aforementioned alternatives. Moreover, an advantage of this method is that the tail dynamics in cyber risk are validated by industry. Given this, it is reasonable to assume that the differences in extreme-event behavior are not so pronounced as to invalidate the use of tail overweighting in this context.

[5] In the next version of the paper, I will continue with this approach and attempt to derive a cdf. However, the remainder of this section already supports the main findings.

[6] This assumption is reasonable. The transformation in equation 3 specifically targets the jump component. Even if the chosen probability measure $\mathbb{P}$ differs from the real $\mathbb{P}$, the estimated risk-neutral measure $\mathbb{Q}$ should exhibit a similar structure compared to the real $\mathbb{Q}$, albeit with a different η. However, since each η inherits the same bias from the $\mathbb{P}$ measure, a structured tail overweighting should still be observable.

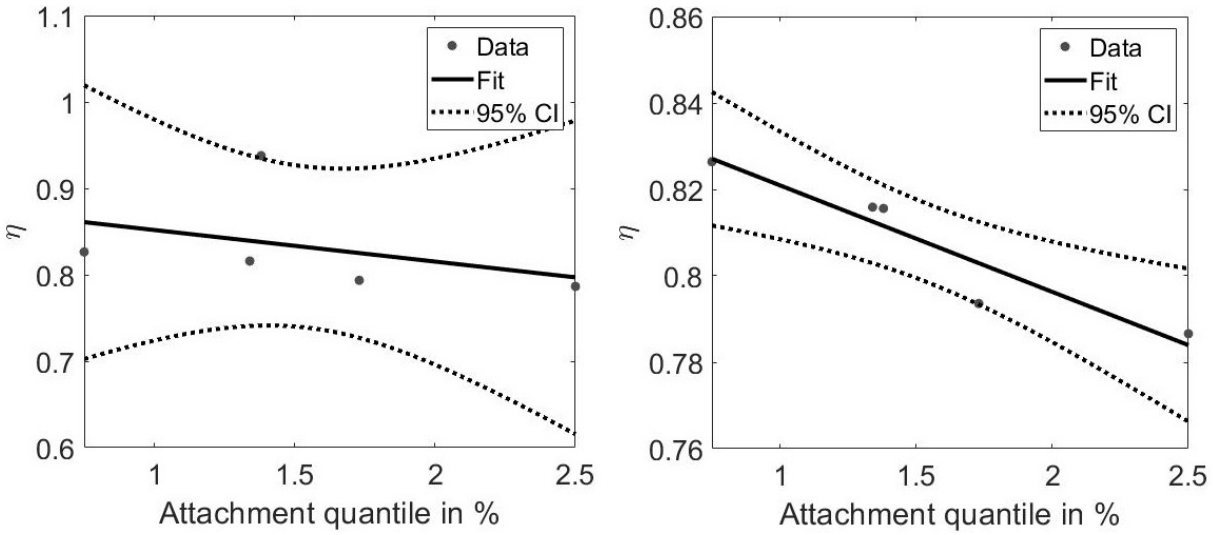
Bond	Date of issue	Modeling agency	Cedent	Attachment point	Expected loss ratio	Spread
1	Feb 2024	AIR worldwide	Tower Hill Insurance Exchange	1.73%	1.56%	10.25%
2	Apr 2024	AIR worldwide	Florida Peninsula Insurance	2.5%	1.78%	9.5%
3	Apr 2024	AIR worldwide	American Coastal Insurance Company	1.38%	0.89%	10.25%
4	May 2024	AIR worldwide	Citizens Property Insurance	1.34%	1.24%	10.5%
5	May 2024	AIR worldwide	Gallatin Point Capital LLC	0.75%	0.65%	7%

Table 7: Public information of the target bonds.

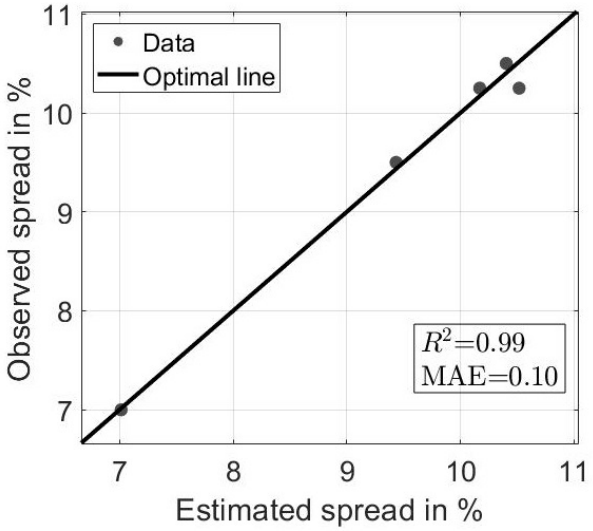
This table presents public information on bonds with the underlying risk peril "Florida named-storm" sourced from Artemis (2025a). For five bonds, it includes the date of issuance, the corresponding modeling agency, and the cedent. Additionally, bond-specific details such as the attachment quantile, the expected loss as a percentage of the notional amount, and the spread as a percentage of the notional amount are also provided.

To keep the appendix short, each individual step is not outlined in detail. The procedure follows the methodology described in section 3. After calibrating the five bonds, the relationship between η and the attachment quantile is examined (Figure 7, part (a)). An R^2 value of 0.15 generally indicates a poor fit. However, a closer look reveals that four bonds align relatively linearly, with only bond 3 from Table 7 deviating as an outlier. This follows a familiar pattern, suggesting that the spread of bond 3 likely incorporates additional risk premiums. To account for this, a beta factor of 3% is assumed to the spread, and the fit is re-evaluated (Figure 7, part (b)). Incorporating the beta factor increases R^2 to 0.89, revealing the expected pattern of η in relation to the attachment quantile. The model from part (b) is then used to estimate the spreads. Figure 7, part (c), demonstrates that the combination of the beta factor and tail overweighting produces a near-perfect fit.

Although the dataset for bond calibration is problematic, the analysis suggests that the observed relationships are not coincidental. A similar pattern of tail overweighting appears in Florida named-storm bonds, reinforcing the structural similarities in downside risk. This finding points to potential direction for further research into the broader applicability of tail-risk dynamics.



(a) Relationship with η w/o beta factor; $R^2 = 0.15$. (b) Relationship with η with beta factor; $R^2 = 0.89$.



(c) Model prediction.

Figure 7: Relationships with η.

This figure presents calibration of the attachment quantile to η (a) without beta factors and (b) with beta factors. The data are in points, the solid line is the regression fit, and the dashed line is the 95% confidence internal (CI). Part (c) is the spread estimation with the beta factor. Furthermore, it shows the R^2 and the mean absolute error (MAE).