Caterina

A Siena Capital Group Company

Siena Capital

A Valuation Methodology for Catastrophe Bonds Incorporating Primary and Secondary Market Activity with Seasonality-Normalized Hazard Rate Dynamics

Prepared by

Jack Stone
Ridwan Bari
Mathys Baldacchino

August 22nd, 2025

1 Introduction

The catastrophe bond (cat bond) market represents a distinct segment of the insurance-linked securities (ILS) sector, with approximately USD 57 billion of risk capital outstanding as of 2025. [1] Cat bonds are issued by insurers, reinsurers, or governmental risk pools to transfer catastrophe risk, such as hurricanes, earthquakes, typhoons, severe convective storms, and certain non-natural perils like cyber into capital markets.

1.0.1 Regional Peril Coverage

While global in scope, the market is concentrated in a few key peril regions:

- **North America** – Primarily U.S. hurricane and earthquake exposure, representing the largest share of outstanding cat bonds. There are various sub-regions within North America that exhibit unique pricing behavior, (i.e., Florida wind risk typically commands a higher expected return than Northeast wind risk).
- **Asia-Pacific** – Japan typhoon and earthquake transactions are the most prevalent, with occasional coverage for other regional perils.
- **Europe** – European windstorm, hail, and earthquake risks.
- **Multi-Region / Aggregate Programs** – Portfolios combining multiple perils and geographies in a single issuance.

1.0.2 Investor Base

The end investor community comprises pension funds, endowments, sovereign wealth funds, and increasingly, high net-worth individuals. Specialist ILS managers dominate allocations, though broader institutional manager participation has grown as the asset class has matured. Many investors seek diversification benefits, as catastrophe risk is largely uncorrelated with broader financial markets.

1.0.3 Market Development

Modern cat bonds emerged in the mid-1990s following Hurricane Andrew (1992) and the Northridge earthquake (1994), when traditional reinsurance capacity proved insufficient. Standardized legal structures and the adoption of 144A issuance enabled scalability and secondary trading. Over the past three decades, the market has grown from niche experiments to a critical source of global reinsurance capacity.

Today, cat bonds provide indemnity, index, parametric or modeled loss-based payouts to cedants, while offering investors clearly defined risk-return profiles and more liquidity compared to traditional reinsurance allocation strategies, though two-way liquidity in the market has recently been limited thanks to investor supply of capital exceeding cedants' demand for protection. This combination of liquidity, diversification, and defined term has made cat bonds an increasingly important allocation for portfolios seeking alternative non-correlated returns.

2 Current Pricing Cadence and Methodologies

2.0.1 Pricing Frequency

Weekly pricing is standard in the cat bond market. The persistence of weekly pricing in the cat bond market is less a function of asset class fundamentals and more a reflection of technological inertia. While other fixed income sectors have moved to near-real-time pricing, cat bond valuations remain anchored to workflows and systems that were designed decades ago. Some data vendors provide "daily pricing", which in reality is a daily refresh of weekly broker price sheets with more recent TRACE prices displayed where applicable. In cases where no such price is available, fund valuation committees have discretion to mark the securities. [2]

2.0.2 Current Price Production Methodologies

A single TRACE trade in a given peril/maturity bucket can imply a repricing for dozens of comparable bonds. Today, this repricing process is performed manually by some ILS specialist broker-dealers. Analysts enter the observed trade into a spreadsheet, observe the implied discount margin (DM), identify the comparable bonds, and make the appropriate DM and price adjustments by hand. This creates latency and the opportunity for inconsistency in how new market data flows into comparable bonds that are not trading.

Some brokers have integrated Python scripts into these spreadsheets to partially automate steps such as baseline comparable analysis fitting or spread recalculation. However, these workflows still depend on generic financial datasets from third-party technology providers that cannot natively model the complexities of catastrophe bonds. This leads to misstated cash-flows, which distort derived yields, discount margins, and ultimately the broker marks that feed the market. Furthermore, various datasets such as loss reports, reset reports, and collateral reports are uploaded throughout the duration of the bond to at least three different deal room providers, requiring brokers and managers to track, ingest, structure, and apply this data into their valuations. To date, there has been no standardization of this data nor its method of delivery, depriving the cat bond market of a unified data schema, streaming in real time, that would be expected of a supposedly liquid and mature capital market.

2.0.3 Potential Aversions to Standardization Among Market Participants

Some specialist managers have long recognized the systematic inefficiency and latency of collecting accurate cat bond data and the resultant errors in broker pricing that arise from the systemic failure of third party vendors to capture and model this bond financial data. In some cases, they have incorporated these pitfalls into trading strategies or even capital-raising narratives, for example, arguing their ability to interpret and clean flawed data is a competitive advantage. Yet no broker nor specialist manager has created an independent, technology-driven, market-wide pricing standard, allowing outdated methods to persist unchallenged. Some specialist managers who have prudently established elaborate internal systems for handling these inconsistencies may be less inclined to support market standardization, recognizing the inefficiencies of the asset class keep large asset managers with superior distribution networks from competing for client market share.

Because broker pricing and data extraction is resource-intensive under the current setup setup, weekly valuation remains the norm. Furthermore, the lack of technological innovation around pricing and valuation supports the narrative that specialist managers are needed to prudently manage cat bond portfolios. (Some specialist managers argue that their maintenance of broker and sponsor relationships, their interpretation of flawed catastrophe models, and their trading of market cycles are other areas where they add value. The authors of this paper believe that under current conditions these arguments obtain, and such arguments will be addressed and engaged with in forthcoming whitepapers.) Producing daily marks and standardizing data for the entire cat bond market is seen as operationally impractical without a full technology rebuild.

3 Detailed Overview of Pricing Methodology

3.0.1 Risk Bucketing Framework

Risk Buckets To enable consistent and defensible repricing, all outstanding cat bonds are assigned to **risk buckets**. These buckets are defined by characteristics that drive comparable risk-return behavior:

1. **Peril** – e.g., U.S. hurricane, Japan typhoon, U.S. earthquake, European windstorm, cyber.
2. **Trigger Type** – i.e., occurrence v. aggregate; indemnity, index, parametric, modelled loss
3. **Risk Period / Remaining Term** – short-dated vs. long-dated bonds experience different seasonality decay profiles.
4. **Loss Development Status** – clean bonds versus loss-impacted bonds with open event exposure.

All bonds within a bucket share a common relative value framework. If one bond in the bucket trades, the new market-implied risk premium is applied across the group, adjusted for each bond's idiosyncratic spread differential.

3.0.2 Sharpe Ratio-Based Auto-Adjustment

Sharpe Ratio Preservation In this framework, the Sharpe ratio is calculated directly from the floater's discount margin and expected loss:

$$SR = \frac{DM_t - EL_t}{\sigma_t}$$

Where:

- SR = target Sharpe ratio at inception
- DM_t = clean discount margin at time t for the floater
- EL_t = seasonality-adjusted expected loss at time t
- σ_t = seasonality-adjusted volatility at time t

The EL_t term is calculated by seasonally adjusting the expected loss of the note as reported in the offering circular or as of the latest applicable reset report. In the case of a multi-peril bond, the percent contribution to expected loss by peril should be applied to the total expected loss, and the expected loss of each peril within the bond should be seasonally-adjusted for the given valuation date. The σ_t term is based on a proprietary formula developed to capture peril-specific volatility decay through the risk period, and it is also seasonally-adjusted. It is used by some hedge funds to price catastrophe risk and reflects the changing probability density of loss as time to peak peril season changes.

3.0.3 Daily Pricing in the Absence of a Market Event

If no market event (TRACE trade or primary issuance) occurs for a given bucket on day t:

1. Hold SR constant at the inception value.
2. Update σ_t for the given valuation date.
3. Update EL_t for the given valuation date.
4. Solve for DM_t from the Sharpe equation:

$$DM_t = SR \cdot \sigma_t + EL_t$$

Apply the updated DM_t to the bond's cashflow model to determine the daily indicative clean price.

3.0.4 Daily Pricing During Market Events

When a TRACE trade or primary issue occurs within a risk bucket, the traded bond(s) establish a *percent change in Sharpe ratio* that is used to update the non-traded bonds in that bucket.

Step 1 — Clean price and compute Sharpe for traded bonds For each traded bond i, the TRACE-reported *dirty price* $P_{t,i}^{\text{dirty}}$ is first *cleaned* to remove accrued collateral interest. Collateral accrual may be linked to a floating reference rate (e.g., SOFR, EURIBOR) or to a managed treasury fund (e.g., JPMorgan, Federated Hermes, Goldman Sachs). The clean price $P_{t,i}^{\text{clean}}$ reflects the bond's market valuation *net of collateral*, isolating the risk spread. This cleaning process is performed because it is cat bond market standard for secondary trade prices to be quoted clean of risk spread and dirty of collateral yield. This standard may change as market participants begin adopting new technology that is able to account for accrued interest on the collateral.

From the clean price, a discount margin $DM_t^{(i)}$ is solved, and the updated Sharpe ratio is calculated:

$$SR_{t,i}^{\text{new}} = \frac{DM_t^{(i)} - EL_t^{(i)}}{\sigma_t^{(i)}}$$

The percent change in Sharpe is then:

$$\text{PctChange}_i = \frac{SR_{t,i}^{\text{new}} - SR_{t,i}^{\text{old}}}{SR_{t,i}^{\text{old}}}$$

Step 2 — Average Sharpe movement if multiple trades If more than one bond trades in the same bucket on the same day:

$$\text{PctChange}_{\text{anchor}} = \frac{1}{N_{\text{traded}}} \sum_{i=1}^{N_{\text{traded}}} \text{PctChange}_i$$

Step 3 — Assign EL tiers All non-traded bonds in the bucket are classified into three tiers based on current expected loss:

- **Low EL:** $0\% \leq EL < 1.5\%$ — least dampened (largest share of the anchor move)
- **Medium EL:** $1.5\% \leq EL < 3\%$ — moderate dampening
- **High EL:** $EL \geq 3\%$ — most dampened (smallest share of the anchor move)

Each tier has a proprietary dampening factor λ_{tier} ($0 < \lambda \leq 1$), with $\lambda_{\text{low}} > \lambda_{\text{med}} > \lambda_{\text{high}}$.

Step 4 — Apply dampened percent change to non-traded bonds For each non-traded bond j in the bucket:

$$SR_t^{(j)} = SR_{t-}^{(j)} \times \left[1 + \lambda_{\text{tier}(j)} \cdot \text{PctChange}_{\text{anchor}}\right]$$

Step 5 — Convert updated Sharpe to DM

$$DM_t^{(j)} = SR_t^{(j)} \cdot \sigma_t^{(j)} + EL_t^{(j)}$$

Step 6 — Convert DM to clean price The updated discount margin $DM_t^{(j)}$ is inverted through the bond's cashflow model to solve for the *clean price*:

$$P_{t,j}^{\text{clean}} = f^{-1}\left(DM_t^{(j)}, \text{Cashflows}_j\right)$$

Here, f^{-1} represents the inverse yield/discount margin function applied to the bond's precise cashflow schedule, accounting for structural features (e.g., pre-risk/post-risk coupons). The resulting clean price can then be converted to a dirty price by re-adding accrued interest, either risk interest, collateral interest, or both.

Key Points:

- Cleaning the TRACE price ensures the movement reflects *only* changes in risk spread, not collateral accrual.
- Adjustment is based on *percent change* in Sharpe ratio, preserving proportional moves across EL tiers.
- Low EL bonds respond most to the anchor change, high EL bonds the least.
- Proprietary dampening factors control cross-bond influence while maintaining bucket integrity.

4 Additional Considerations

4.0.1 Data Integrity

Within a strict mark-to-market framework, the incorporation of all observable secondary market transactions is essential for regulatory defensibility and methodological consistency. Accordingly, every TRACE-reported transaction is incorporated into the daily valuation process, irrespective of trade size, counterparty relationship, or apparent transaction motivation (e.g., portfolio transfers, affiliate trades). Excluding trades on the basis of subjective assessments would introduce discretion into the process, undermining both auditability and alignment with fair value accounting principles. All TRACE prices are normalized to remove collateral accrual effects prior to their use in Sharpe ratio calculations.

4.0.2 Latency

Valuations are produced on a trade-date basis, with all available TRACE observations for the session incorporated into end-of-day marks. This approach ensures that the daily mark reflects the complete set of contemporaneous market observations while maintaining procedural consistency. A technological platform which incorporated the aforementioned methodology with an event-driven architecture would be able to easily increase calculation frequency from daily, to intraday, to live pricing. The existence of such a platform would potentially increase liquidity by attracting large market making firms and bringing further standardization and professionalization to the asset class.

4.0.3 Volatility Controls

No explicit cap is applied to the magnitude of Sharpe ratio changes derived from traded instruments. Although this approach allows for material movements in category-level pricing when anomalous trades occur, it preserves the integrity of observed market signals and avoids the introduction of human-induced bias. Transmission of such movements to non-traded instruments within the same risk bucket is moderated solely through pre-defined dampening factors, which are a function of the instrument's expected loss tier. While these dampening mechanics objectively reduce volatility, their design is informed by direct observations of how the secondary market behaves and reprices risk following a trade. In practice, the magnitude of repricing across similar instruments is a function of their expected loss profile, rather than a deliberate effort to engineer synthetic price stability.

4.0.4 Methodological Philosophy

The methodology is principally concerned with the accurate reflection of observable market conditions rather than the artificial suppression of volatility. While smoothing techniques can create the appearance of price stability, such interventions generate synthetic signals that may distort true market dynamics. By contrast, the present framework preserves the informational content of all observed trades, even when doing so results in some day-to-day variability in reported marks, which would be expected in any liquid capital market.

References

[1] Steve Evans. Catastrophe bond issuance breaks annual record already in 2025 at over $17.8 bn. *Artemis.bm*, July 2025. Accessed: 2025-08-20.

[2] Stone Ridge Asset Management LLC. Stone ridge interval fund prospectus. Prospectus (PDF), March 2025. Accessed: 2025-08-20.