A New Model for Weathering Risk: CDOs For Natural Catastrophes

Aaron C. Koch, FCAS, MAAA

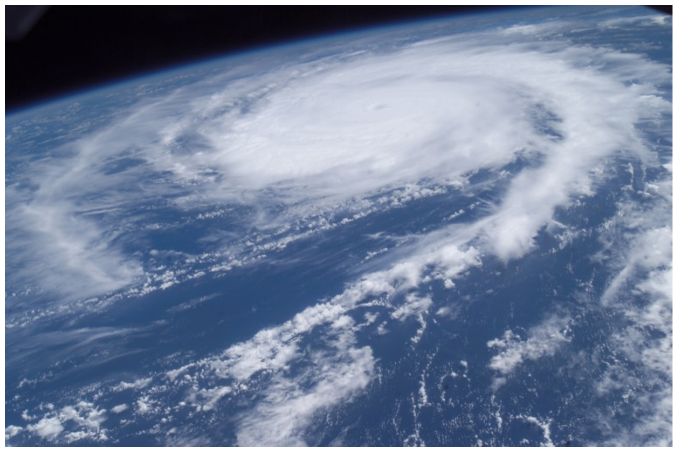

Table of Contents

1. INTRODUCTION

A fundamental function of reinsurance is to provide financial recovery from natural catastrophes. A well-functioning insurance market can enable rebuilding efforts that would overwhelm the resources of individual households and communities.

By its nature, catastrophe risk is often not diversifiable on a local or even regional scale. Instead, insurers usually look to the global reinsurance markets for catastrophe risk protection.

Recently, a third option has risen to prominence as hedge funds, pension funds, and other institutional investors (hereafter "alternative capital") have sought to "directly" invest in catastrophe risk. Through investments in insurance-linked securities (ILS) and collateralized reinsurance, alternative capital has increased the available supply of property catastrophe risk coverage and driven ILS prices toward all-time lows.

Yet alternative capital may be most remarkable not for its impacts to date, but for its vast untapped potential. ILS is still a niche asset class, and the existing market for catastrophe risk is dwarfed by the pool of available institutional capital. Nevertheless, price competition on existing property catastrophe risk may have already reached the point of diminishing returns. Product innovation is needed to support the growth rate of alternative capital and to produce further improvements in the availability and cost of catastrophe coverage.

In this paper, the use of pooling and tranching techniques similar to those used in collateralized debt obligations (CDOs) is proposed as a tool for expanding the market for catastrophe risk.

Given the somewhat infamous legacy of CDOs and inherent complexity of catastrophe risk, the pairing of the two may seem problematic. However, catastrophe risk is a far stronger candidate for inclusion in a CDO-type structure than economic assets such as securitized mortgages. An appropriately designed *collateralized risk obligation* (CRO) would have significantly less systemic vulnerability than the subprime mortgage-fueled CDOs at the heart of the recent financial crisis.

In fact, the concept of a CRO is not entirely novel. Limited numbers of CRO-type instruments were issued in the early to mid-2000s, only to largely disappear at the onset of the financial crisis. The market for catastrophe risk has matured in the interim, yet suffers from limitations that a CRO is well-suited to address.

CROs should improve efficiency and stimulate growth in the catastrophe risk market in two key ways.

First, CROs would *simultaneously increase the availability of investment-grade catastrophe risk and high-yielding catastrophe risk*. The vast majority of recently securitized catastrophe risk is either unrated or assigned a speculative, or "junk," rating. Large institutional investors frequently place limits on the amount of non-investment-grade risk they will hold in their portfolios. On the other end of the spectrum, falling yields on catastrophe risk have led to heightened demand for high-yielding catastrophe securities. Thus, CROs should further expand the supply of alternative capital.

Second, CROs should *encourage investment in heretofore underinsured perils and geographic locations*. In a CRO, unusual or diversifying perils provide enhanced value, which is due to their low correlation with the other assets in the portfolio. CROs should contribute to the globalization of a predominantly US and European market, and stimulate the growth of insurance in developing economies.

Guide to this paper

Building a CRO: Sections 2 and 3

Section 2 presents a brief overview of alternative capital, its recent increase in popularity, and the emerging need for product innovations such as the CRO.

Section 3 outlines a basic design framework for CROs. It also addresses a key question: How are CROs different from the CDOs that underpinned the financial crisis? This section demonstrates that the primary pitfalls of pre-crisis CDOs are largely mitigated for CROs, which is due to the nature of insurance risk and the structure of insurance markets.

Case study: Section 4

Section 4 provides a stylized example of a CRO. This "sample CRO" is used to discuss potential pricing and rating methodologies for CROs. In addition, it illustrates the differences in achievable credit enhancement (i.e., leverage) between CROs and traditional CDOs.

Practical considerations, market history, and conclusion: Sections 5 and 6

Section 5 considers several practical and historical questions surrounding the implementation of a CRO. Who are the likely sponsors? What kinds of CRO-type instruments were issued prior to the financial crisis? What lessons can be drawn from their history?

Section 6 provides conclusions.

2. THE RISE OF ALTERNATIVE CAPITAL

2.1 – Market transformation

2013 was a banner year for alternative capital investments in their various forms (see Sidebar 1). Catastrophe bonds enjoyed their second largest issuance year on record, and reached an all-time peak for the amount of total principal outstanding (approximately US$20.2 billion).[1] Collateralized reinsurance had even stronger growth, surpassing the traditional catastrophe bond market in size for the first time.[2] This significant influx of alternative capital led to rapidly falling prices, with some sources quoting a year-over-year decrease in catastrophe bond spreads of nearly 40%.[3]

Yet, this may be just the beginning of alternative capital's entry into catastrophe risk markets. While current estimates peg the amount of invested alternative capital at around US$50 billion as of early 2014, this pales in comparison to the estimated US$30 *trillion* of existing worldwide pension fund assets. Further allocations to catastrophe risk of even 1% to 2% of these assets could double or triple the capacity of the existing US$300 billion USD catastrophe risk market.[4] Many market analysts expect an explosive next five years for alternative capital, with projections ranging from $40 billion USD to $150 billion USD in *new* alternative capital entering the marketplace.[5,6,7]

However, for every risk investor there must also be a risk seller, and there are limits on alternative capital's ability to grow purely through price competition on property catastrophe risk. If supply

Sidebar 1:
Major Existing Types of Alternative Capital Investments

Catastrophe bonds: Investments in Special Purpose Vehicles (SPVs) in which a limit of catastrophe coverage is fully collateralized by outside investors, who are paid periodic risk-based coupons by the ceding (re)insurer through the SPV. After a triggering event, the investor may lose the principal for the ceding reinsurer to cover claims.

Collateralized reinsurance: Reinsurance coverage in which capital markets investors fully collateralize the reinsurance limit offered in exchange for an up-front reinsurance premium.

Industry Loss Warranties (ILWs): Dual-trigger reinsurance or derivative contracts (typically fully collateralized) in which the payout is based upon both an industry loss threshold and the ultimate net loss to the cedent.

Sidecars: Financial structures designed to allow outside investors to take on a quota-share portion of the risk written by a (re)insurer, by establishing a collateralized limit of coverage for which reinsurance premiums are paid. Generally designed to have a limited lifespan and intended to capture the increase in rates often witnessed after a major catastrophe.

---

[1] Swiss Re [33]
[2] Artemis.bm [2]
[3] Plenum Insurance Linked Capital [26]
[4] Guy Carpenter [14]
[5] Artemis.bm [3]
[6] Artemis.bm [4]
[7] BNY Mellon [8]

outpaces demand, investors will at some point reach a minimum acceptable return for a given level of risk. Indeed, some reports have suggested that prices on certain risks have already begun to reach this lower boundary.[8,9] Alternative capital also faces a stiff test from traditional providers of catastrophe risk protection. Despite public promises to avoid a pricing "race to the bottom," catastrophe-focused reinsurers are unlikely to simply let profitable business walk away. Longtime client relationships and add-on services provide reinsurers an edge that in some instances may overcome the lower prices of alternative capital.

Thus, the long-term growth prospects for alternative capital depend on the ability to leverage its primary advantage over the traditional reinsurance model—a lower cost of capital—into the development of *market-expanding* and *market-completing* innovations.

2.2 – Market expansion

Market-expanding innovations introduce new exposures to the alternative capital market. To a certain extent, the globalization of the insurance industry will sow the seeds of opportunity for *market-expanding* innovation. As the epicenter of insurance growth shifts toward Asia-Pacific and similar regions,[10] opportunities for investing in new catastrophe risks will multiply. The quality of catastrophe models and data for these regions will also improve, providing potential investors with better tools for measuring risk and assessing investment opportunities.

Alternative capital investors have already demonstrated enthusiasm for the limited number of developing market catastrophe securitizations to date. Catastrophe bonds for diversifying perils—such as Mexican hurricane risk, Mexican earthquake risk, and Turkish earthquake risk—have enjoyed high investor demand and coupon spreads well below the market average. More opportunities may be on the horizon, as officials in countries such as India and the Philippines have recently voiced interest in securitizing a portion of their countries' catastrophe risks.[11,12]

Another candidate for *market-expanding* innovation is the securitization of new types of risk, including terrorism risk and catastrophic liability risk. These risks are significantly harder to model than natural catastrophe risk. For instance, any terrorism risk model must contend with terrorists' intention of avoiding predictability. Nevertheless, falling margins in property catastrophe risk may eventually push alternative capital into these harder-to-model perils: At some point, every risk must have its price.

[8] Munich Re [24]
[9] Carrier Management [11]
[10] Munich Re [23]
[11] The Economic Times [13]
[12] Reuters [29]

2.3 – Market completion

Market-completing innovations bring new investors into the market. Despite recent success, the catastrophe risk market still lacks many of the characteristics exhibited by more mature financial markets. Certain *market-completing* innovations stand to grow the catastrophe risk market by removing existing supply-side limitations.

For example, most catastrophe bonds to date[13] have been identified as speculative or "junk"-grade risk, receiving ratings between BB- and BB+. Compared to other fixed-income alternatives such as corporate bonds, catastrophe bonds are disproportionately high-risk investments.

The low ratings of existing catastrophe risk instruments serve to limit the pool of potential alternative capital investors, as institutional investors are frequently limited in the amount of non-investment-grade risk they can hold. Similarly, financial services companies subject to risk-based capital standards (e.g., banks and insurers) are required to hold more capital for low-rated, non-investment-grade assets.

2.4 – The role of the CRO

Assume that you wish to build a new product for the property catastrophe risk market with both *market-expanding* and *market-completing* properties. These goals would appear to be at odds. Market expansion usually requires the inclusion of previously uncharted risks, entailing greater uncertainty and requiring correspondingly higher returns. Conversely, market completion through the introduction of investment-grade investment options seems to require the creation of *safer*, lower-risk catastrophe instruments.

The introduction of CROs may be able to achieve both of these objectives. CROs would produce a wide spectrum of rated catastrophe risk, opening up investment opportunities for a broader range of potential investors. It would also promote the securitization of new types of risk in order to help diversify the catastrophe risk assets collateralizing the pool.

Of course, catastrophe risk is very different from the credit risk found in traditional CDOs. As such, it is important to understand how a CRO might differ from a traditional CDO, and the impacts of these differences on the success and sustainability of the CRO.

[13] The current market also contains a number of unrated instruments offering investment opportunities generally similar to speculative rated securities.

3. DIVIDE AND CONQUER: CATASTROPHE RISK AND THE CRO

3.1 –CDOs and insurance companies

Set aside catastrophe risk for a moment and begin with the underlying structure: How does a basic CDO function? During the financial crisis, CDOs (particularly those containing high concentrations of risky subprime mortgages) became notorious for their complexity and inscrutability. However, these derivative structures bear striking structural similarities to a much older and more familiar type of financial vehicle—the insurance company. The resemblances are illustrated by two key parallels between CDOs and insurers: tranching and the law of large numbers.

A CDO is comprised of a pool of financial assets[14] carved into *tranches,* a series of ordered claims to the pool's cash flows. Investors in *senior tranches* have first claim to pool profits, and are followed in order by holders of *mezzanine* and *equity* (or *junior*) *tranches*. In exchange for bearing a larger share of the pool's default risk, equity trancheholders are compensated with the highest potential returns. At the opposite end, senior tranches appeal to risk-averse investors willing to accept lower returns in exchange for holding highly rated assets.

The appeal of a CDO is that the most senior tranches[15] can often be structured to satisfy rating agencies' requirements for an exceptionally strong (typically AAA) credit rating. Generally, the credit ratings of these tranches significantly exceed those of the underlying pool collateral assets were they to be rated individually, which is due to the security provided by the subordinate tranches.

The structure of an insurance company is fundamentally similar. As with a CDO, an insurer carves up a pool of underlying assets (in this case, the profits or losses on insurance policies) into *de facto* "tranches." The tranched structure of an insurer is illustrated by considering the priority order of the insurer's liabilities in a run-off scenario. Outstanding policyholder obligations (e.g., loss and unearned premium reserves) receive the highest priority, and are analogous to senior CDO tranches. The insurer's other debt obligations are equivalent to mezzanine tranches, and the equityholders of an insurer match to equity CDO trancheholders (see Figure 1). This tranche-based description of insurers has been examined in detail elsewhere, notably in the context of analyzing reinsurance arrangements.[16]

---

[14] These assets are often assumed to be homogenous.
[15] Complex CDO structures can have upwards of 10 (or more) tranches, with several senior tranches, a number of *mezzanine* tranches, and one or more equity tranches.
[16] Mango, D. & Bunick, C. [21]

Figure 1: Structural Comparison of a CDO and an Insurance Company

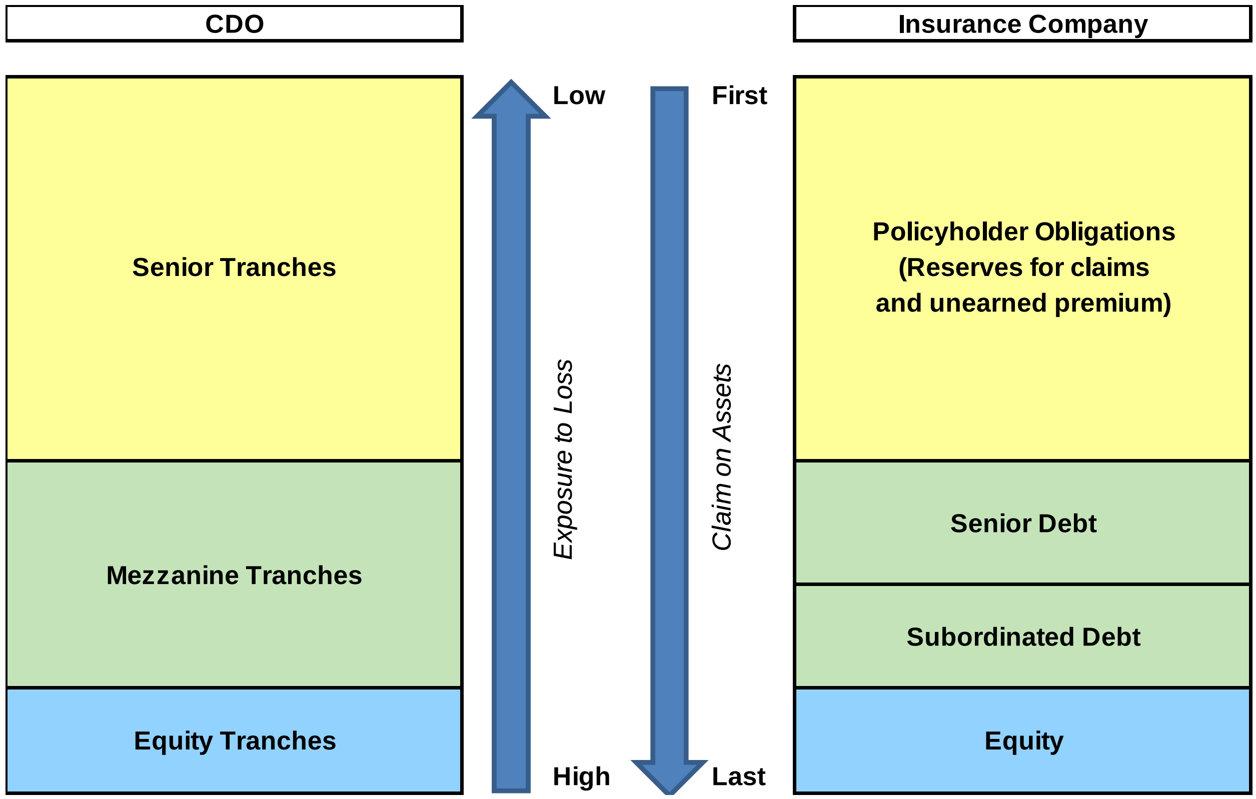

The second parallel is that CDOs, like insurers, derive their economic value from the law of large numbers. Both are vulnerable when the assumptions underlying the that law does not hold and single events can cause a large volume of highly correlated losses—Hurricane Andrew and the mortgage downturn being prominent examples.

These similarities beget an obvious question from comparing many decades of insurance industry success to the financial conflagration caused almost immediately by CDOs: Why the enormous discrepancy? Further, could a CRO avoid the pitfalls that undermined the CDOs of the mid-2000s?

The remainder of this section seeks to answer these questions. First, the composition of a CRO will be roughly outlined and compared to that of a pre-crisis CDO. Then, the CRO will be scrutinized in the context of the key structural factors contributing to the collapse of the CDO market.

3.2 – Designing a CRO

As noted above, catastrophe risk can pose a threat to the law of large numbers. By nature, catastrophes affect a large number of policies simultaneously, making catastrophe risk diversifiable only on a global scale.

There is, however, potential for diversification - of two kinds - if sufficient variety of global catastrophe risk can be collected into a single pool. First, this structure would enjoy *geographic* diversification: For instance, Asian typhoon activity is not fully correlated (and in fact may be inversely correlated) with North American hurricane activity.[17]

Second, a broadly based catastrophe risk pool benefits from *typological* diversification. Natural disasters can be geophysical (e.g., earthquake), meteorological (e.g., convective storms and tropical cyclones), or even climatological (e.g., drought) in nature.[18] Each type is driven by forces that are not fully correlated—and often, are not correlated at all. For example, the occurrence of a Japanese earthquake is unlikely to impact the likelihood of a major U.S. hurricane.

The ideal CRO will pool the broadest range of natural catastrophe risks possible to ensure ample diversification. But what form will these risks take? The most basic building block of insurance risk is the individual insurance policy, but securitized instruments are ill-suited to insuring single policies. Instead, the pooled "risks" in a CRO must be pooled *portfolios* of catastrophe risk collected by insurers.

Sidebar 2:
Impact of Asset Quantity on a CDO

High (1,000+ assets):

Benefits

- Achieves greater spread of diversifiable risk
- Creates proportionally more highly-rated securities
- Allows for greater structuring flexibility and complexity

Drawbacks

- Requires significantly stronger modeling assumptions
- Creates more of a "black box" – complexity may not fully be understood
- Requires higher ongoing management costs

Low (5-10 assets):

Benefits

- Simpler to model – may be able to fully specify relationships between each pair of assets
- Uses less resources for gathering and managing pool assets

Drawbacks

- Achieves less credit enhancement
- May be harder to reach cost-efficient pool size

The cost efficiency of the CRO is further enhanced if the pooled catastrophe risks are already securitized and tradable. Catastrophe bonds have a somewhat liquid secondary market, and have already undergone the initial modeling and pricing process. Fractional shares of existing catastrophe bonds likely represent strong building blocks for the CRO.

But does the market have enough existing catastrophe risk material to support CROs? Outstanding catastrophe bonds number in the dozens, while a single mortgage-based CDO pooled thousands of individual mortgages. It was this numerousness (and the law of large numbers) that enabled the credit enhancement found in CDOs.

Nonetheless, it should not be necessary to acquire thousands of catastrophe assets to create a

---

[17] Maloney, E. & Hartmann, D. [20]

[18] Among other types not listed.

well-functioning CRO. The key to the feasibility of the CRO is the fact that a relatively low number of pooled assets is needed to capture the vast majority of portfolio diversification benefits. For instance, under the assumptions shown in Figure 2, approximately 85% of the available diversification benefits are captured by a pool of only 10 securities. After this point, the marginal benefit of adding further assets to the pool decreases rapidly.

Figure 2: Diversification Effects on Portfolios of Varying Size

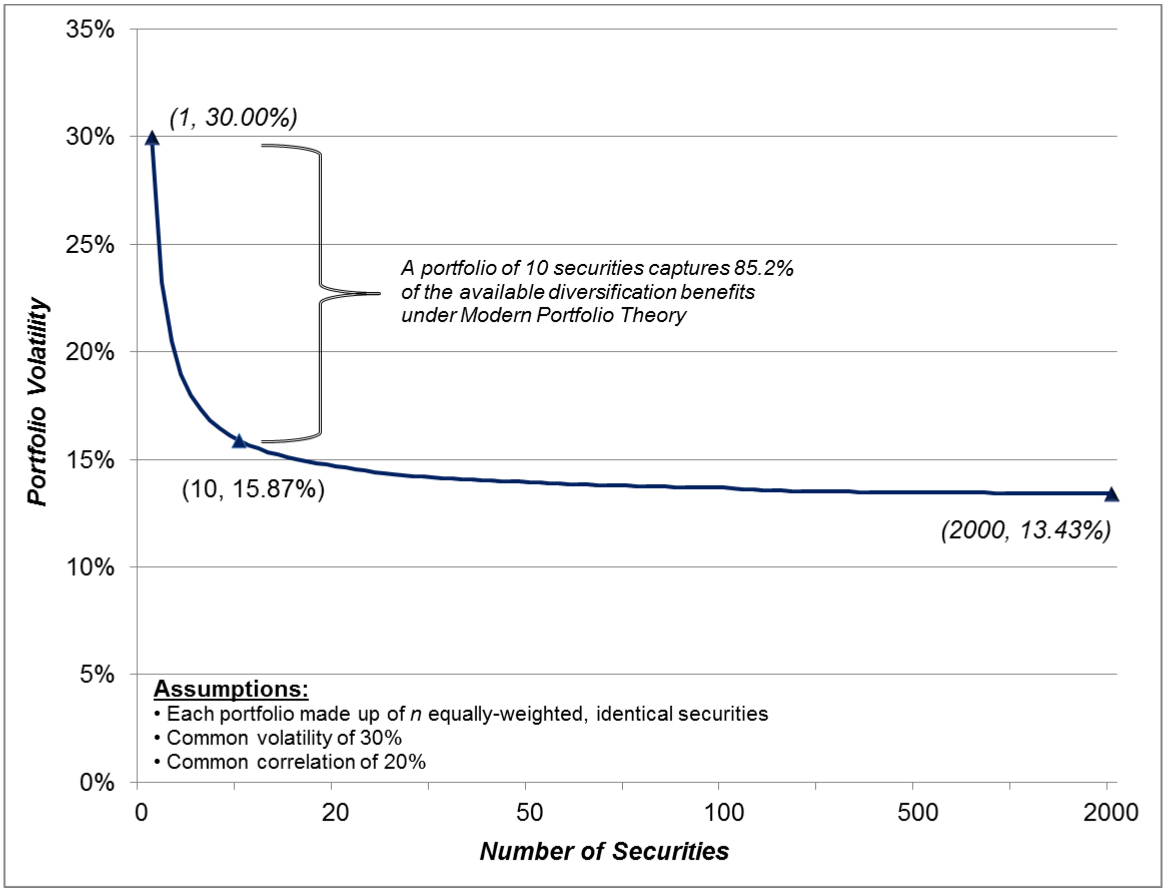

Compared to catastrophe risk securities, individual mortgages are comparatively low in value: A large number of mortgages are required to create an economically viable pool.[19] On the other hand, individual catastrophe bonds are frequently issued with several hundred million dollars of principal at stake, each bond covering a portfolio of thousands of individual insurance policies. A fractional share of a single catastrophe bond offering can represent an investment of many millions of dollars.

It should be feasible to create a CRO with a relatively low number of underlying assets, perhaps between five and 10. The resulting structure will likely be smaller than the average pre-crisis subprime mortgage CDO (for which one source provides an average size of $829 million).[20]

[19] Ashcraft, A. and Schuermann, T. [6]
[20] Barnett-Hart, A.K. [7]

However, CROs should also have a lower size threshold for economic viability, because of the reduced expense load for ongoing management of several securities as opposed to thousands.

3.3 – Does the CRO have the same vulnerabilities as pre-crisis subprime CDOs?

The effort spent designing a CRO is wasted if the resulting structure exhibits the same weaknesses that led to the collapse of the CDO market during the financial crisis. What are the key factors that led to those losses, and how should we expect a CRO to fare in comparison?

After the onset of the crisis, many sought to diagnose the causes behind the collapse of the CDO market. Their conclusions, while wide-ranging, tended to highlight similar themes. These themes can be separated into two categories:

- **Modeling-focused observations:** Addressed *which* assumptions failed to match reality.
- **Behavioral-focused observations:** Addressed *why* assumptions failed to match reality.

The balance of this section provides a brief overview of the fall of pre-crisis CDOs through the lens of the categories above. It finds that a well-designed CRO should fare better on almost every test of systemic vulnerability, proving to be significantly more robust than its pre-crisis subprime predecessors.

3.4 – Modeling: The actuary and the Gaussian copula

In standard CDO models, two types of input parameters must be estimated for each underlying asset. The simpler is the *default profile*—the likelihood of default across time, independent of any other asset in the pool. If this information is not readily available for each asset, then simplifying homogeneity assumptions can streamline the model.

A more significant challenge is quantifying an asset's *dependency profile* with each of the other assets—in short, how its likelihood of default is affected by surrounding defaults. For a portfolio of 1,000 mortgages, modeling the interactions among each asset using a traditional linear correlation matrix requires close to a half million parameter estimates.[21] This approach is usually unwieldy—and for a long time represented the biggest barrier to CDO modeling.[22]

In 2000, actuary David X. Li proposed pricing CDOs with copula models, which were then primarily found in biostatistics and actuarial science.[23] In particular, Li's paper presented the use of a Gaussian copula constructed from a multivariate normal distribution. Unlike many other copula forms, the Gaussian has the practical advantage of being easily generalizable from the two-variable

---

[21] Because the correlation coefficient is assumed to be 1 along the diagonal, and is the same pairwise for analogous cells on either side of the diagonal.

[22] Mathematically, it is nearly impossible to ensure a positive semi-definite correlation matrix.

[23] Li, D. [18]

situation (find the single dependency between *x* and *y*) to the *n*-variable situation (find the various dependencies among *a*, *b*, *c*, *d*…etc.). To do so, it utilizes a crucial simplification: It assumes that the entire dependency structure for the pool of assets is driven by a common factor, which can be estimated as a single pool default correlation parameter ρ (rho).

The Gaussian copula model (and the similar models that followed) provided solutions to what had seemed an impossible mathematical problem—but in return, it required the assumption that the relationships among thousands of mortgages could be fully expressed by a solitary constant. This key parameter held enormous sway over the model output, particularly due to the significant leverage inherent in CDOs. Thus, the most senior CDO tranches were almost indestructible *assuming the models used to price and rate them were correct*—and highly susceptible to downgrade if they were not.[24,25,26,27]

The models, of course, turned out to be wildly optimistic. As the housing bubble burst, mortgage default rates skyrocketed past all recent historical benchmarks, nationwide. This caused rating agencies to reassess their models and downgrade AAA tranches at an unprecedented pace. The ensuing collateral calls and liquidity crunch kicked off the financial crisis and crystallized the public's image of the CDO: A structured finance vehicle both incomprehensible and toxic.

In comparison, CROs should be able to rely upon a more accurate and robust modeling process. Natural catastrophe models forecast physical events, while economic models forecast human behavior: The former lie more in the realm of science, the latter social science. While it is important not to downplay the amount of uncertainty in catastrophe models (which is significant), it is also true that they need not capture the additional behavioral component inherent in financial markets - which often drives tail outcomes (e.g., a "run on the bank"). At least in the short run, humans can have little impact over the occurrence or severity of any particular natural catastrophe—the very reason that catastrophe risk is desired as a *zero-beta* investment.[28]

Further, the low number of assets in a CRO allows for more transparent pricing. With accurate exposure information and access to catastrophe models, it is possible to research each CRO asset in detail.[29] In a CRO, relationships among specific assets can be identified and modeled on a case-by-case basis as opposed to relying on a single, catch-all assumption to represent the entire dependency structure. Potential CRO modeling techniques are described in Section 4.2 below.

---

24 Coval, J., Jurek, J., & Stafford, E. [12]
25 Krahnen, J. & Wilde, C. [17]
26 Hull, J. & White, A. [16]
27 Heitfield, E. [15]
28 That is, an investment showing no correlation with the performance of the equity markets as a whole.
29 Some CDOs, to further complicate the matter, were comprised solely out of tranches of other CDOs – creating a third layer of tranching.

3.5—Behavior: A matter of incentives

While it is important to understand the weaknesses in the CDO pricing models used in practice during the mid-2000s, it is more important to understand *why* the modeling assumptions turned out to be wholly inaccurate.

At the heart of the matter was an incentives problem: Most of the key participants in the life cycle of a CDO were compensated according to the volume of completed CDO transactions. Worse, most CDO originators retained little to none of the downside risk associated with the securities:

- Mortgage writers adopted an "originate-to-distribute" model that removed their portfolios of subprime mortgages from their balance sheets and led to a loosening of loan standards.[30]
- Major banks then turned loan portfolios into securitized instruments, collecting a healthy underwriting fee while typically retaining little to no risk.[31] The lower-rated tranches of these mortgage-backed securities were then re-tranched into CDOs, providing yet another opportunity for fees.[32]

A misguided incentive structure plagued not only the formation of CDOs, but their evaluation by the major credit rating agencies. As the CDO market exploded, so too did the fees paid to rating agencies - who were paid not only on volume, but by the *arranger* of the security (and not by the ultimate investor). For the ratings agencies, taking a more pessimistic view than the competition often meant watching arrangers take their subsequent (and highly lucrative) business elsewhere.[33]

As a result, the only participants with a strong incentive to accurately assess the quality of the assets were the investors themselves. In reality, many investors were either unable or unwilling to invest the resources necessary to obtain their own view of CDO risk, instead putting their faith in the CDOs' sterling credit ratings. Only the eventual market implosion revealed what is obvious in hindsight: Real skin in the game—fundamental to appropriately motivating CDO intermediaries—was absent at nearly every stage.

Fortunately, these issues are largely absent in a CRO. Securitized insurance risk is generally written on an *excess-of-loss* basis, often with cedent co-participation in the reinsured layer. Unlike pre-crisis mortgage originators, primary insurers thus have every incentive to write good business—because they retain the vast majority of risk on their policies.

In addition, reinsurance markets are highly relationship-based. Many of the same intermediaries

---

[30] Purnanandam, A. [28]
[31] Ashcraft Schuermann [6]
[32] Many of the losses eventually suffered by major banks were incurred by those that weren't quite *good enough* at removing the toxic assets they were creating from their balance sheets.
[33] Barnett-Hart [7]

that serve the catastrophe bond market also provide services for a range of non-catastrophe reinsurance transactions by similar (or even the same) parties. There arguably exists a stronger reputational incentive for the actors along the catastrophe securitization structuring chain to respect the interests of the other involved parties.

Thus, a properly structured CRO may avoid both the modeling and incentives problems that plagued pre-crisis CDOs. Significant expertise is still needed to grasp the numerous sources of risk inherent in catastrophe contracts: Nevertheless, the robustness of insurance markets should allow the CRO to avoid becoming simply the latest example of a "toxic" structured finance asset.

4. A CRO PRICING EXAMPLE

4.1 – Overview

This section considers the pricing of a theoretical five-asset CRO. It is designed to highlight the key features of the structure without excessive functional detail. Analysis is presented in a simplified form, with most mathematical details left to the Appendix.

For our sample CRO, we assume that the asset pool consists of fractional shares of single-peril Rule 144(A)—that is, publicly issued—catastrophe bonds with a one-year duration. Each bond has thus been evaluated by a third-party catastrophe model vendor during the initial pricing and issuance process. We assume that we have a stand-alone exceedance probability (EP) curve for each asset representing its loss profile. We also have the following summary statistics for each bond:

- **Attachment probability:** The likelihood that a bond will suffer a nonzero loss to its principal.
- **Expected loss:** The average percent of principal that a bond is expected to lose.
- **Exhaustion probability:** The likelihood that a bond will suffer a complete loss to its principal.

Figure 3 shows the selected parameters for each of the five securities in our sample CRO, based loosely on existing market securities. The current, public Standard & Poor's rating table for catastrophe bonds was utilized to estimate a rating for each security. None of the securities listed below qualifies for an investment-grade credit rating.

Figure 3: Sample CRO Composition

	Term (Years)	Assumed Expected Loss	Estimated Attachment Probability	Estimated Exhaustion Probability	S&P Implied Rating
Florida Hurricane - FLH	1	4.00%	5.33%	2.67%	B+
New England Hurricane - NEH	1	2.00%	2.67%	1.33%	BB
US (California) Earthquake - USQ	1	3.00%	4.00%	2.00%	BB-
Japan Earthquake - JPQ	1	3.00%	4.00%	2.00%	BB-
Turkey Earthquake - TUQ	1	1.50%	2.00%	1.00%	BB+

Next, we aim to estimate similar performance metrics for each CRO tranche.[34] To do so, we must first define the tranches' attachment and detachment (or exhaustion) points.

In a CDO, tranches are designed to maximize the size of the highest-rated (usually AAA) tranches. The endpoints of each tranche are frequently calibrated to precisely meet the minimum standards of a given rating category. With only five assets, such a level of refinement is less feasible for a CRO. Instead, we will divide the sample pool into four illustrative tranches:

- **Equity tranche:** Eroded by aggregate losses of the first 0% to 20% of pool collateral.
- **Mezzanine tranche:** Eroded by aggregate losses from 20% to 40% of pool collateral.
- **Senior tranche:** Eroded by aggregate losses from 40% to 60% of pool collateral.
- **Super-Senior tranche:** Eroded by aggregate losses from 60% to 100% of pool collateral.

4.2 – CRO dependency modeling

The prices of CRO tranches, as with those of a traditional CDO, depend heavily on how the individual risks in the collateral pool relate to one another. These effects are magnified greatly in senior tranches, whose loss profiles are highly leveraged on the pool's dependency patterns.

There are at least two plausible approaches to reflecting asset dependencies in CRO pricing.

The first approach is through *event simulation*, using the same techniques that are used to price stand-alone catastrophe bonds. Each major catastrophe modeler produces a simulated "event set" of natural catastrophes. Using this event set, simulated years are generated for the portfolio of exposures—which in this case represents all of the securities within the asset pool. From these simulations, one can obtain the EP curve and summary statistics referenced above. These are then used to price the CRO tranches.

However, investors may consider the outputs of the major catastrophe modelers to be somewhat of a black box. Further, perhaps detailed exposure information is not available for each asset in the CRO, or the model becomes unwieldy and hard to analyze on a portfolio-wide basis. In any of these cases, investors may desire another approach for establishing their own "view of risk."

An alternative to event simulation is a *portfolio analysis* approach. For our sample CRO, this consists of taking each of the five individual EP curves (i.e., catastrophe model outputs) and relating them to one another—as opposed to trying to create a comprehensive portfolio EP curve out of the combined exposure sets.

[34] Note that for a CRO with sequentially stacked, non-overlapping tranches, the exhaustion probability of one tranche is equivalent to the attachment probability of the next higher tranche.

Catastrophe modeling firms already enable this type of analysis by producing pairwise linear correlation matrices for the existing universe of catastrophe bonds.[35] However, using this particular type of matrix for CRO pricing may be problematic. It best suits analysis based on the normal distribution—the use of which was one of the issues with traditional CDO pricing.

Recall that normal Gaussian copulas (and related forms) were used for CDO pricing because of the challenges in extending two-dimensional dependency modeling to higher dimensions. Fortunately, there are alternative methodologies that are likely superior for pools with relatively few assets, such as a CRO.

The use of *vine copulas* is one such alternative. Vine copulas tackle the multi-dimensional challenge by modeling pairs of copulas in two dimensions and then linking them together in "vines." Importantly, this procedure allows for the use of tail-heavy copulas that fit the tail-heavy distributions being modeled—and better, it allows for the use of *different* copulas to model each pair of assets. This eliminates the need for an overarching (and often ill-fitting) assumption about which single copula best fits the data.

4.3 – Rating the sample CRO with a vine copula approach

A vine copula model can be used to estimate the loss parameters and credit ratings of the various tranches of our sample CRO.[36] To illustrate the importance of asset interdependencies, we rate the CRO tranches under two assumptions:

- *Independence model:* The performance of each asset in the CRO is assumed to be fully independent of the performance of each of the other assets.
- *Dependencies model:* While independence is assumed for many pairs of assets, a few are assumed to have positive loss correlation captured by a *Clayton copula*:[37]
  - There is assumed to be a slight positive correlation between Florida and Northeast hurricane risk.
  - Similarly, there are assumed to be slight positive correlations among earthquake risks in the U.S., Japan, and Turkey.
  - There is assumed to be no correlation between hurricane and earthquake risk types.

Figures 4 and 5 show the results of modeling under each set of assumptions. Modeling details, mathematical derivations, and additional key assumptions are outlined in the Appendix.

---

35 Risk Management Solutions, Inc. Miu Platform [30]

36 After the key loss statistics and ratings have been estimated, the tranches can be priced through any number of theoretical approaches or by comparison to existing market securities.

37 For the purposes of this paper, the various *taus* are selected judgmentally and for illustrative purposes only.

Figure 4: Estimated Independence Model Tranche Loss Parameters

Sample CRO Independence Model				
Tranche	Tranche Range	Default Probability	Expected Loss	Implied Rating
Junior	0-20%	16.79700%	12.86343%	CCC+
Mezzanine	20-40%	1.03720%	0.65358%	BB+
Senior	40-60%	0.02670%	0.01458%	A
Super-Senior	60-100%	0.00020%	0.00003%	AAA

Based on Monte Carlo Simulation

Based on the Independence model, we find that both the senior and super-senior tranches of the sample CRO have simulated default probabilities low enough to quality for investment-grade ratings on the S&P ratings table for structured finance instruments.[38] In addition, the super-senior tranche (representing a full 40% of the CDO collateral) meets the ratings table standards for a AAA rating.

Figure 5: Estimated Dependencies Model Tranche Loss Parameters

Sample CRO Dependencies Model				
Tranche	Tranche Range	Default Probability	Expected Loss	Implied Rating
Junior	0-20%	16.62980%	12.75911%	CCC+
Mezzanine	20-40%	1.17590%	0.74806%	BB+
Senior	40-60%	0.04010%	0.02176%	A-
Super-Senior	60-100%	0.00050%	0.00006%	AA+

Based on Monte Carlo Simulation

In contrast, the Dependencies model concentrates a higher percentage of the losses in tail outcome events—that is, in higher tranches of the CRO. As a result, both the senior and super-senior tranches of the sample CRO are rated one notch lower than the comparable tranches in the Independence model.

[38]Barnett-Hart [7]

Even though the absolute effects on default probability and expected loss may be small (e.g., the expected default probability on the super-senior tranche goes up 0.0003%), the *relative* impacts of the Dependencies model are significant, as shown in Figure 6.

Figure 6: Modeled Loss Comparisons by Tranche

Sample CRO Model Comparison				
	(1)	(2)	(3)	(4)
	Independence Model Expected	Dependencies Model Expected	(2) - (1)	(3) / (1) Relative Percent
Tranche	Loss	Loss	Difference	Change
Junior	12.86343%	12.75911%	-0.10431%	-0.811%
Mezzanine	0.65358%	0.74806%	0.09448%	14.456%
Senior	0.01458%	0.02176%	0.00717%	49.164%
Super-Senior	0.00003%	0.00006%	0.00003%	73.380%

Based on Figures 4 and 5

The incorporation of dependencies shifts risk from the junior tranche to the other tranches, with the relative impact growing as the level of seniority increases. This finding, which matches the results of a number of financial crisis analyses, highlights the importance of accurate incorporation of dependencies into a CRO model—despite the fact that correlations among insurance risks may be low compared to those found in the financial markets.

4.4 – Other observations

Despite the potential advantages a CRO may enjoy in terms of modeling accuracy and stability, it nevertheless provides less credit enhancement (e.g., fewer AAA-rated assets) than a traditional CDO. Figure 7 compares our sample CRO tranching structure to that of a pre-crisis mortgage-backed CDO—Goldman Sachs' GASMP Trust 2006-NC2 (GSAMP). GSAMP was split into the following 17 tranches:[39]

- 5 senior tranches, each rated AAA by S&P
- 9 mezzanine tranches, with investment-grade ratings ranging from AA+ to BBB-
- 3 junior tranches (including an equity tranche), with speculative-grade ratings ranging from BB+ to unrated.

[39] Ashcraft Schuermann [6]

Figure 7: Tranche Structure, Pre-crisis CDO vs. Sample CRO

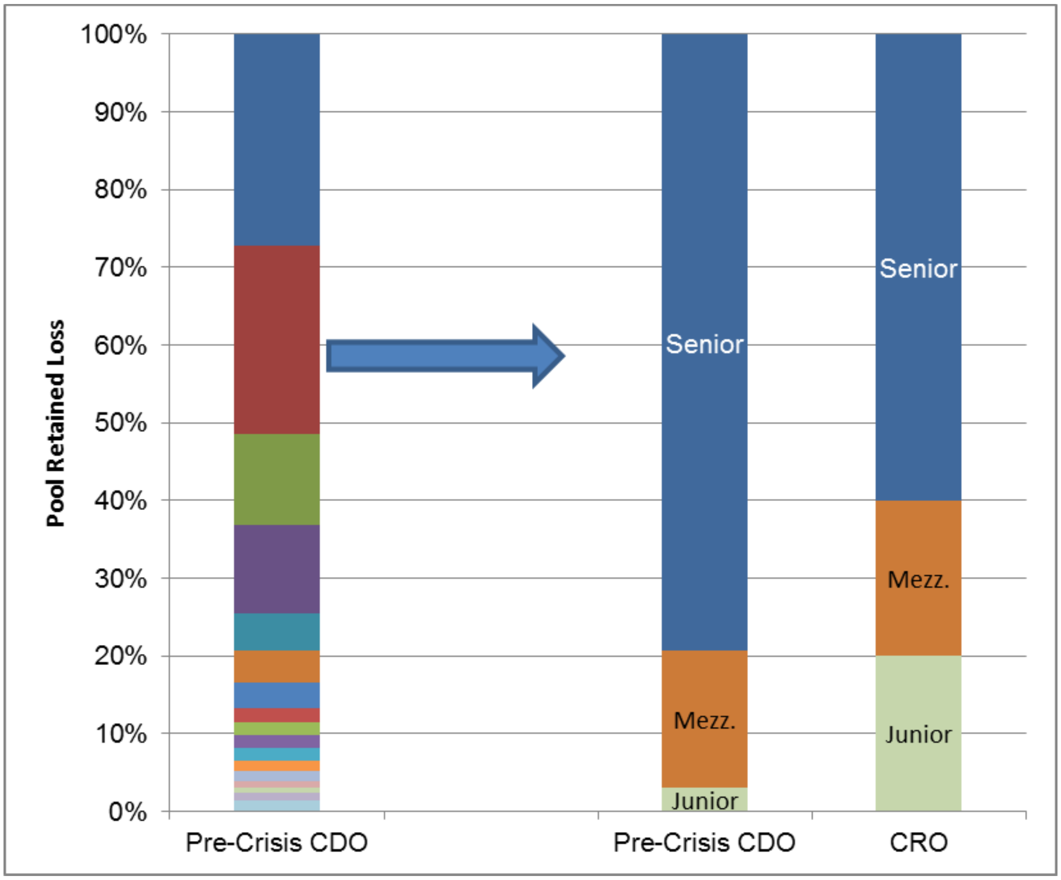

GSAMP contains a significantly higher proportion of senior risk than the sample CRO. On a proportional basis, the difference is most striking between the sets of junior tranches, which make up 20% of the CRO but only 3% of GSAMP.

However, this may not be too severe of a drawback for the CRO. In some instances, the alternative capital market has shown significant appetite for ILS with a default risk similar to or higher than that of the sample CRO's junior tranche. For example, USAA's Residential Re 2013-2 (Class 1) catastrophe bond exposed investors to an attachment probability of 21.38% and an expected loss of 13.06% - and received one of the lowest pricing ratios of coupon-to-expected loss ever seen in the market.

In fact, the CRO's creation of junk junior tranches may be a significant benefit to alternative capital investors who seek a high-yielding portfolio of catastrophe risk as opposed to a well-diversified one.[40] For these investors, rapidly falling spreads on peak perils such as Florida hurricane

[40] The rationale for these investors is often that holding catastrophe risk itself (in small quantities) serves as the macro-level diversifier for the rest of their portfolios. Under this paradigm, pursuing diversification within the catastrophe portfolio can lead to an unnecessary erosion of returns.

risk have threatened their return objectives. This potentially opens the door for the use of small, targeted allocations to CRO junior tranches as part of a larger investment strategy.

5. POTENTIAL SPONSORS AND CAPITAL CONSIDERATIONS

5.1 – Market history: Multiple-event securitizations

Prior to the financial crisis, a few major insurers and reinsurers experimented with creating investment-grade securitizations of their risk. These *multiple-event securitizations* worked similarly to CROs: High credit ratings were obtained by insuring only the second, third, or further subsequent events happening in a given period across a worldwide portfolio. In essence, they were CROs that simply excluded the junior tranche.

The first catastrophe bond to have a tranche receive an "A" rating in such a manner was issued by the French reinsurer SCOR Group in December 2001.[41] Atlas Reinsurance II covered European windstorm, Japanese earthquake, and Californian earthquake risk. It had two tranches: Class B notes provided coverage for the second qualifying catastrophe in the contract period, while Class A notes provided coverage for the third.[42] While the Class B notes received a BB+ rating, Class A received the coveted "A"—with an annual expected loss of 0.05% and coupon spread above LIBOR of 2.38%.[43,44]

This rating reversed S&P's policy of maintaining a BBB+ ceiling for catastrophe bonds, which was due to the "cliff risk" inherent in a first-event cover: No matter how unlikely the event, the owner of a first-event catastrophe security faces the risk of full default with little or no warning. Because Atlas II required an accumulation of events to be triggered, S&P was comfortable that cliff risk was sufficiently mitigated. In the event of a first triggering event, investors would have the opportunity to reassess their holdings - and offload them if they believed the risk profile no longer suited their objectives.[45]

Other major insurance players such as Swiss Re and Converium also issued multiple-event securitizations in the early to mid-2000s. Since then, the market for such products appears to have largely disappeared: The share of investment-grade catastrophe risk fell from roughly a quarter of the market in 2007 to zero in 2013.[46] Why might this have occurred?

---

41 Woo, G. [37]
42 SCOR retains the first qualifying catastrophe without assistance from Atlas Reinsurance II.
43 Woo [37]
44 Artemis Deal Directory [5]
45 Boyle C. [9]
46 Willis Re Capital Markets [36]

5.2 – Capital adequacy: The challenge of single-cedent CROs

To illustrate, consider the *return periods* (the "1 in X" odds) for triggering each tranche of the sample CRO under the Dependencies model. By implied rating, the mezzanine and senior tranche are approximately equivalent to Atlas II's Class B and A notes, respectively.

Figure 8: Return Periods by Tranche, Sample CRO

Sample CRO Dependencies Model		
Tranche	**Default Probability**	**Return Period**
Junior	16.62980%	6 Years
Mezzanine	1.17590%	85 Years
Senior	0.04010%	2,494 Years
Super-Senior	0.00050%	200,000 Years

The senior tranche of the sample CRO has a return period of approximately 2,500 years. In comparison, the Solvency Capital Requirement (SCR) under Solvency II requires European insurance companies to hold capital to protect against a 1-in-200-year series of events. Most companies hold capital to protect against events beyond this 1-in-200-year standard: Ratings agencies will generally insist on it. However, holding capital is expensive. Companies may not care (and it could be argued, *shouldn't* care) about risk at return periods far exceeding their internal capital adequacy targets.

A similar line of argument provides a case against a single-cedent CRO. For example, a company with the risks contained in the sample CRO and a capital adequacy horizon that doesn't extend to 2,500 years could simply buy protection up to the exhaustion point of the mezzanine tranche (e.g., by securitizing only the first two events occurring on the global portfolio). This would prevent the company from paying for coverage that is not in line with its overall strategic plan.

There are still a number of reasons for a company to consider a full-fledged CRO solution. Perhaps an insurance group's regionally based companies need reinsurance cover that cannot be consumed by catastrophes in another region, or the company has capital management goals that go beyond a simple analysis of return period adequacy. Nevertheless, it is not surprising that the popularity of the single-cedent multiple-event securitization has waned over time.

5.3 – Market history: Gamut Re

An alternative vision of the CRO (and the primary one offered in this paper) combines securitized risks from a number of companies. This avoids the problems described in the prior section, as each company's risks can be securitized at a lower return period (e.g., the 25- to 100-year periods commonly found in today's catastrophe bonds). For this type of CRO, the historical precedents are more infrequent—perhaps limited to a single structure established immediately prior to the financial crisis.

In June 2007, the hedge fund Nephila Capital raised over $300 million to sponsor Gamut Re Ltd., a sidecar-type vehicle whose returns from investing in catastrophe risk were allocated across five tranches.[47,48] The catastrophe portfolio held in Gamut was actively managed by Nephila Capital, and ran through the end of 2009. Details are shown in Figure 9

Figure 9: Gamut Re Tranche Structure.

Nephila Capital - Gamut Re Tranche Structure, Ratings, and Yields			
Class	Size (in M $USD)	Coupon*	S&P Rating
A	60	1.4%	A-
B	120	3.0%	BBB+
C	60	7.0%	BB-
D	25	15.0%	NR
E	45	Equity	NR

*Represents spread over LIBOR

Source: PR Newswire

Gamut Re expired at the end of 2009 and was not renewed, with Nephila citing the increased cost of debt in the immediate post-crisis markets.[49] Since then, it appears that no similar transactions have been attempted (at least publicly).

Yet conditions have changed since 2009. Fixed-income coupon rates currently approach all-time lows and an unprecedented (and increasing) number of investors have turned their attention to the catastrophe risk markets. As the diversity and number of catastrophe securitizations continue to

[47] PR Newswire [27]
[48]SIFMA: Insurance & Risk-Linked Securities Conference [31]
[49] Trading Risk [34]

increase, so too does the feasibility of the CRO—and with it, the expansion of catastrophe risk markets to investors and risks that heretofore have remained on the outside looking in.

6. CONCLUSION: NEW SKIES AHEAD?

More than most other economic assets, catastrophe risk securitizations are well-suited to inclusion in tranche-based leveraging structures. As evidenced by the lessons of the recent financial crisis, such structures do not offer a panacea for maturing financial markets: Nevertheless, the CRO may serve as a powerful tool for completing and expanding the existing market for catastrophe risk.

Because of their inherent similarity to CDOs—the fuel for the financial crisis meltdown—we can expect that the concept of tranched catastrophe risk might require patient exploration. However, a well-structured CRO is likely to avoid many of the systemic modeling and incentive-based vulnerabilities that were fatal to the pre-crisis CDO market. In contrast, CROs are well-positioned to take advantage of recent advances in dependency and catastrophe modeling to provide a nuanced, powerful, and relatively transparent basis for market analysis.

Above all, the CRO is arguably the optimal tool for generating investment-grade catastrophe risk, a missing ingredient in the current market. Securitizing risk that is too far out in the tail (either on a single-event or multiple-event basis) is unlikely to appeal to many companies on a stand-alone basis. As the catastrophe risk market continues to expand, however, it becomes increasingly possible to generate investment-grade risk by combining the risk of a number of different companies—opening up new possibilities for the financial markets to spread the risk from natural catastrophes on a global basis.

APPENDIX: MODELS FOR CRO PRICING

A.1 – Key modeling assumptions and background

The following assumptions and background discussion apply to each of the two models discussed in this Appendix.

Assume that the assets are equally weighted in the CRO—that is, each asset represents 20% of the overall portfolio. Further, assume that the catastrophe bonds are zero-coupon: That is, a single payment is made at the end of the payment term to the bondholder if the bond has yet to default.[50]

On an individual basis, we assume that the payoff profile for each catastrophe bond asset in the pool could be replicated by a bull put spread (assuming the appropriate options existed). For example, a catastrophe bond written to cover a 50% quota share of the layer $100 million excess of $100 million would have a payoff profile as a function of *X*, the random variable representing size of loss, as shown in Figure 10.

Figure 10: Sample Catastrophe Bond Payoffs by Size of Loss

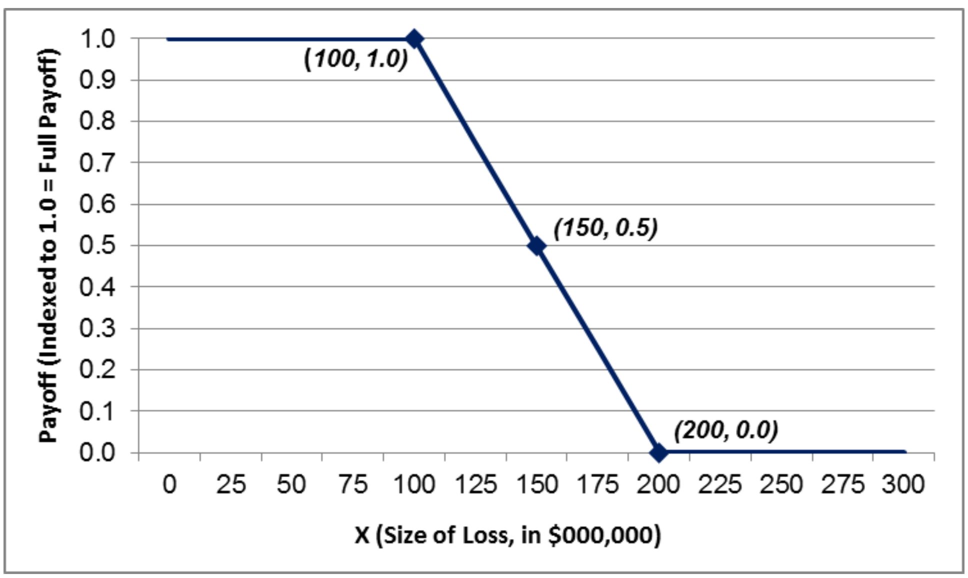

By combining the payoff function with the exceedance probability function (that is, the survival distribution function of the size of loss), we can derive the full payoff probability function for each asset. In practice, the exceedance probability function will be given by the modeled EP curve. For the sake of this example, we assume that each of our assets has the following simplified payoff structure:

[50] This assumption likely would not be justified in an actual pricing model, but does not materially change the generalizable conclusions of this paper.

- For each asset, assume the chance of loss hitting the insured layer (the "attachment probability") is y_i, which is given by the historical data.
- Similarly, assume that the chance of a full-limits loss to the insured layer (the "detachment probability") is z_i, which is given by the historical data.
- Assume that the exceedance function is linear between the attachment and detachment points.[51] Let the expected loss (EL) for each asset be the probability-weighted expectation of the amount of payoff not received by the bondholder that is due to catastrophe loss. Given this, $EL_i = (y_i + z_i)/2$

For each of the assets in our sample CRO, we assume a term-to-expiration of one year and default parameters designed to approximate current market offerings. We determine the implied credit rating for each asset based on the most recent Standard and Poor's (S&P) rating matrix for catastrophe securities, shown in Figure 11.[52]

Figure 11: Illustrative S&P Catastrophe Risk Rating Table

Portion Of Nat-Cat Risk Factor Table

(%)	aaa	aa+	aa	aa-	a+	a	a-	bbb+	bbb	bbb-
1	0.003	0.010	0.015	0.025	0.040	0.060	0.085	0.234	0.353	0.547
2	0.027	0.048	0.074	0.106	0.150	0.200	0.264	0.514	0.825	1.279
3	0.052	0.085	0.133	0.188	0.260	0.340	0.443	0.850	1.405	2.177
4	0.076	0.123	0.191	0.269	0.370	0.480	0.621	1.246	2.073	3.213
5	0.100	0.160	0.250	0.350	0.480	0.620	0.800	1.704	2.812	4.359
	bb+	**bb**	**bb-**	**b+**	**b**	**b-**	**ccc+**	**ccc**	**ccc-**	
1	1.632	2.525	3.518	4.510	5.824	8.138	23.582	45.560	66.413	
2	3.211	4.946	6.915	8.885	11.751	16.674	38.104	59.145	79.233	
3	4.758	7.230	10.095	12.960	17.152	24.004	46.752	64.835	82.905	
4	6.276	9.380	13.037	16.694	21.921	30.025	52.288	68.078	84.581	
5	7.763	11.403	15.745	20.087	26.089	34.945	56.158	70.313	85.650	

A.2 – The Independence model

The following section prices the CRO tranches under the assumption that the performance of each of the underlying single-peril securities is unrelated to each other asset in the CRO.

The concept and mathematics behind the Independence model are simple. Given the assumption of independence, it becomes a straightforward three-step process to assess the risk profile for each CRO tranche:

---

[51] Outside of the range between the attachment and detachment points, the payoff function is constant.

[52] Standard &Poors [32]

1. Using a Monte Carlo simulation generator, simulate X number of years of performance for each of the assets using the payoff profiles and exceedance curve given above. In this paper, $X = 1{,}000{,}000$.
2. For each simulated year, add up the probability-weighted losses for the assets to get the total loss as a percent of pool collateral.
3. Assign the total pool losses to tranches according to the tranching algorithm outlined in Section 4.

The results are presented in Figure 4 above.

A.3 – The Dependencies model: Clayton copulas

The Dependencies model produces a more nuanced view of the risk profile for each tranche. We take a vine copula modeling approach in conjunction with a set of Clayton copulas. The Clayton copula concentrates risk into the left tail of the dependency structure, which in this case we will take to mean high-loss outcomes leading to low-payout states of the security. A major benefit of the Clayton copula is that it is solvable in closed form, leading to a relatively straightforward simulation process when one of the variables is already known.

This procedure, per Venter (2007), is as follows:[53]

- u and v represent the inverse single-variable cumulative distribution function for x and y respectively, that is:
  - $u = F_X^{-1}(x)$
  - $v = F_Y^{-1}(y)$
- a represents the *Kendall's tau* (τ) for the relationship between the two single-variable distributions. See below for description of Kendall's tau.
- $C_u(u,v)$ represents the partial first derivative with regards to the first argument

Then, the following holds for the Clayton copula (Figure 12):

[53]Venter, G. [35]

Figure 12: Derivation of Simulation Formula for Clayton Copula

$$C(u,v) = [u^{-1/a} + v^{-1/a} - 1]^{-a}, \quad a > 0$$

$$C_u(u,v) = [u^{-1/a} + v^{-1/a} - 1]^{-a-1} u^{-1-1/a}$$

$$\frac{C_u(u,v)}{u^{-1-1/a}} = [(1-u)^{-1/a} + (1-v)^{-1/a} - 1]^{-a-1}$$

$$[\frac{C_u(u,v)}{u^{-1-1/a}}]^{1/(-a-1)} + 1 - u^{-1/a} = v^{-1/a}$$

$$\left\{[\frac{C_u(u,v)}{u^{-1-1/a}}]^{1/(-a-1)} + 1 - u^{-1/a}\right\}^{-a} = v$$

Thus—with knowledge of $C_u(u,v)$, u, and constant a – we can simulate the variable v as the output variable conditioned on the independently simulated variables $C_u(u,v)$ and u.

To utilize this model, we must first estimate Kendall's tau. Although Kendall's tau differs from the standard Pearson product-moment correlation coefficient, its form is much the same—a number between -1 and 1 (inclusive) with the following meanings:

- A tau of -1 represents a pair of fully anti-correlated (negatively correlated) assets
- A tau of 0 represents a pair of uncorrelated assets
- A tau of 1 represents a pair of fully correlated assets

The impacts of Kendall's tau on the modeled outputs from the Clayton copula are shown below. Simulated relationships between variables using a Kendall's tau of 0.0 (Figure 13) and a Kendall's tau of 0.4 (Figure 14) are shown below. Note that while the joint distribution of the variables is evenly spread in Figure 13, it instead shows a concentration in the lower left-hand corner (and a smaller amount of concentration in the upper right-hand corner) in Figure 14.

Figure 13: Clayton copula simulation, Tau = 0.0

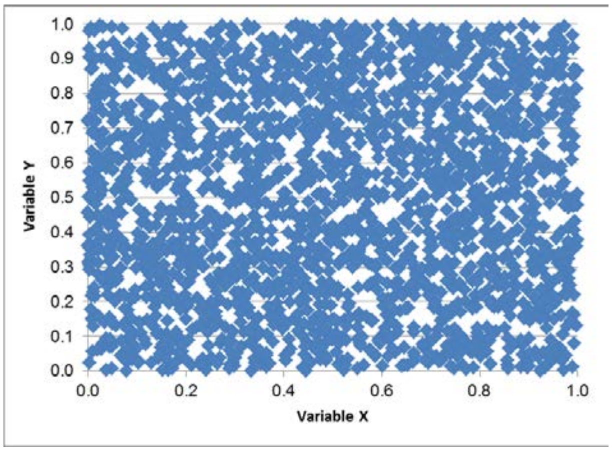

Figure 14: Clayton copula simulation, Tau = 0.4

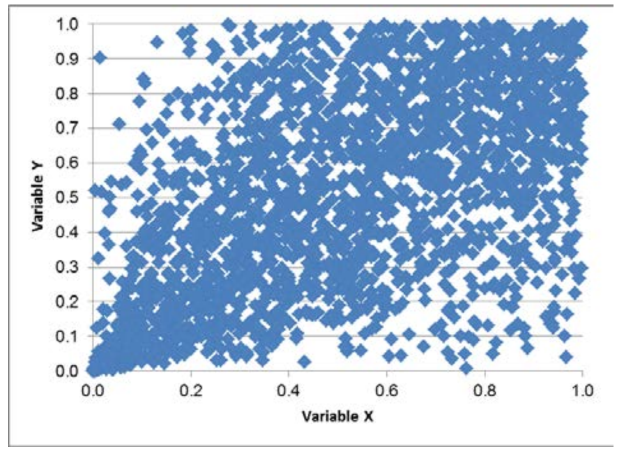

In practice, *tau* estimates will likely be selected for each pair of assets based on a combination of historical data, existing catastrophe models, and expert judgment. For the purposes of this paper, estimates of tau are selected judgmentally and for purely academic purposes as follows:

- The relationship between Florida and New England hurricane risk is captured by a tau of 0.2
- The relationship between US earthquake and Japan earthquake risk is captured by a tau of 0.2
- The relationship between US earthquake and Turkey earthquake risk is captured by a tau of 0.2
- All other relationships (direct or conditional) are independent (captured by a tau of 0.0).

These relationships are captured in the pairwise tau matrix shown in Figure 15 below. Note that despite being captured in a similar form, these constants do *not* represent the linear Pearson's constants typically shown in correlation matrices.

Figure 15: Tau Matrix for Sample CRO

	FLH	NEH	USQ	JPQ	TUQ
FLH	**1.0**	**0.2**	**0**	**0**	**0**
NEH	*0.2*	**1.0**	**0**	**0**	**0**
USQ	*0*	*0*	**1.0**	**0.2**	**0.2**
JPQ	*0*	*0*	*0.2*	**1.0**	**0**
TUQ	*0*	*0*	*0.2*	*0*	**1.0**

A.4 – The Dependencies model: Vine copulas

V*ine copula* models enjoy a number of advantages over the higher-dimensional copulas traditionally used by financial practitioners to model CDOs:

Traditional copula models are limited to the use of a single copula to describe the entire dependency structure. In comparison, a vine copula model is highly flexible: A different copula may be selected for each relationship, reflecting its specific attributes (e.g., tail heaviness).

As a result, vine copula models require far fewer assumptions regarding the behavior of the pool, particularly regarding the homogeneity of pool assets.

The primary multivariate copula models (e.g., Gaussian, Student's t) are generally either symmetric or have only moderate tail heaviness, particularly in higher dimensions. Thus, these models may fail to capture the true tail risk contained in a CDO, particularly given that modeling asymmetries in the dependency structure (i.e., choosing the shape of the copula) can sometimes have a greater impact on results than modeling asymmetries in the marginal distributions of the individual assets themselves.[54]

The primary weakness of a vine copula model is the large number of parameter estimates needed as the size of the pool grows. Assuming a single-parameter copula is used for each dependency, a pool of n assets requires the estimation of (n)(n-1)(0.5) dependency parameters for a fully specified model. The number of estimated parameters can frequently be reduced by careful vine structuring and/or a constant parameter assumption for all conditional dependencies past a certain vine level. Nevertheless, vine copula models are likely to be far more accurate for pools containing a limited number of securities, where the additional precision of the individual dependency estimates is not overwhelmed by the increased risk from estimating many parameters.

Recent empirical testing of financial return data suggests that vine copulas offer improvements over existing models for pools of up to 10 to 12 assets.[55] As a result, CROs are ideal candidates for vine copula modeling, allowing for more precise pricing specifications than previously available for a tranche-based security.

There are a number of ways to build vine copula models, and there are a number of sources offering more detailed explorations of the theory of vine copula modeling.[56,57,58] For the purposes of this paper, a *D-Vine* copula model will be used to evaluate the sample CRO under the Dependencies

[54] Low, R. K. Y. et al. [19]
[55] Ibid.
[56] Nikoloulopoulos, A., Joe, H., & Li, H. [25]
[57]Brechmann, E. & Czado, C. [10]
[58]Mendes, B., Semeraro, M., & Leal, R. [22]

model. For illustration, a D-Vine structure for four assets is shown in Figure 16.[59]

Figure 16: D-Vine Copula Modeling Structure

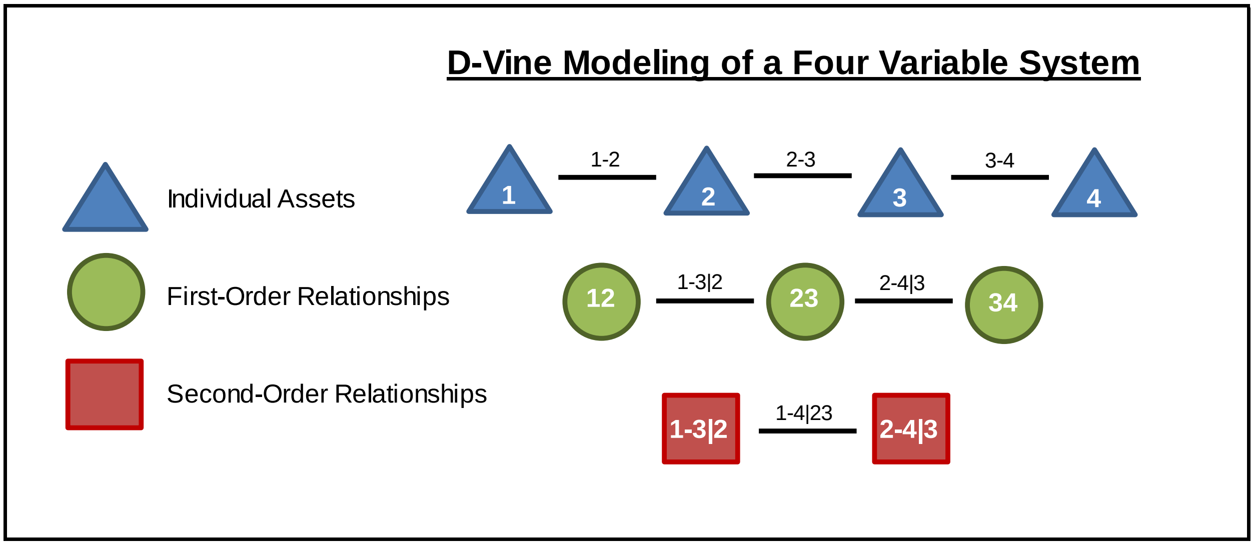

The first level in the vine contains direct pairwise dependencies. Subsequent levels contain *conditional* dependencies based on relationships identified in higher vine levels and conditioned on the shared variables.[60] Once a pool of assets is decomposed into pairwise direct and conditional relationships, pool results are simulated recursively.[61]

With a modified version of the Monte Carlo procedure used above and the simulation methodology provided in Aas et al. (2006),[62] we obtain the revised tranche estimates of default probability and expected loss shown in Section 4.

---

[59] The other prominent vine types are C-Vines and R-Vines, respectively.

[60] For instance, the combination of the 1-2 and 2-3 relationships results in the dependency between variables 1 and 3, conditioned on 2.

[61] Using the inverse of selected partial derivatives of the copula function.

[62] Aas, K. et al. [1]

REFERENCES

[1] Aas, K. et al (2006). Pair-copula constructions of multiple dependence.

[2] Ashcraft, A. and Schuermann, T. (March 2008). Federal Reserve Bank of New York Staff Reports: Understanding the Securitization of Subprime Mortgage Credit.

[3] Artemis.bm (September 9, 2013). Alternative capital growth driven by collateralized reinsurance: Swiss Re. Retrieved March 26, 2015, from http://www.artemis.bm/blog/2013/09/09/alternative-capital-growth-driven-by-collateralised-reinsurance-swiss-re/

[4] Artemis.bm (November 21, 2013). ILS market to double by 2018, drive further rate decreases: Credit Suisse analysts. Retrieved March 26, 2015, from http://www.artemis.bm/blog/2013/11/21/ils-market-to-double-by-2018-drive-further-rate-decreases-credit-suisse-analysts/

[5] Artemis.bm (September 17, 2013). Next $100 billion of alternative capital will transform reinsurance. Retrieved March 26, 2015, from http://www.artemis.bm/blog/2013/09/17/next-100-billion-of-alternative-capital-will-transform-reinsurance-aon-benfield/

[6] Artemis Deal Directory. Retrieved March 26, 2015, from http://www.artemis.bm/deal_directory/atlas-reinsurance-ii-plc/

[7] Barnett-Hart, A.K. (March 19, 2009). The Story of the CDO Market Meltdown: An Empirical Analysis. Harvard College: Cambridge, Massachusetts.

[8] BNY Mellon (November 2013). The Disaster Gap: How Insurers and the Capital Markets can Harness Big Data to Close the Gap.

[9] Boyle C. (January 28, 2002). SCOR Launches Multi-Continental CAT Bond with RMS Technical Help.Insurance Journal. Retrieved March 26, 2015, from http://www.insurancejournal.com/magazines/features/2002/01/28/18787.htm

[10] Brechmann, E. & Czado, C. (April 17, 2013). Risk management with high-dimensional vine copulas: An analysis of the Euro Stoxx 50Plenum Insurance Linked Capital (February 2014). Annual Review 2013: CAT Bond Market and Insurance Events.

[11] Carrier Management (April 27, 2014). Some ILS Money Irrational; Trad'l Re Mkts Adding Terror to Cat Cover: CEO. Retrieved March 26, 2015, from http://www.carriermanagement.com/news/2014/04/27/121784.htm

[12] Coval, J., Jurek, J., & Stafford, E. (April 2008). Economic Catastrophe Bonds.

[13] The Economic Times (May 2, 2014). General Insurance Corporation to offer Catastrophe Bonds. Retrieved March 26, 2015, from http://economictimes.indiatimes.com/markets/bonds/general-insurance-corporation-to-offer-catastrophe-bonds/articleshow/34506135.cms

[14] Guy Carpenter (2014). State of the Market: January 1, 2014 Property Review. RAA Conference 2014.

[15] Hull, J. & White, A. (May 2009). The Risk of Tranches Created From Residential Mortgages.

[16] Heitfield, E. (January 2009). Parameter Uncertainty and the Credit Risk of Collateralized Debt Obligations.

[17] Krahnen, J. & Wilde, C. (April 28, 2008). Risk Transfer With CDOs.

[18] Li, D. (April 2000). On Default Correlation: A Copula Function Approach.

[19] Low, R. K. Y. et al. Canonical vine copulas in the context of modern portfolio management: Are they worth it?

[20] Maloney, E. & Hartmann, D. (2001).The Madden-Julian Oscillation, Barotropic Dynamics, and North Pacific Tropical Cyclone Formation. Part I: Observations. Journal of the Atmospheric Sciences.

[21] Mango, D. & Bunick, C. (2012).Reinsurance Evaluation Using Capital Tranching.CAS Spring Meeting 2012. Retrieved March 26, 2015, from https://cas.confex.com/cas/spring12/webprogram/Handout/Session5036/Mango_Bunick%20%20Tranching%20Capital%20and%20Reinsurance%20CAS%202012.pdf

[22] Mendes, B., Semeraro, M., & Leal, R. Pair-Copulas Modeling in Finance.

[23] Munich Re. Insurance-Linked Securities (ILS) Market Update Q1 2014

[24] Munich Re (March 20, 2013). Asia-Pacific will be the dominant driver of world insurance growth by 2020. Retrieved March 26, 2015, from http://www.munichre.com/en/media-relations/publications/press-releases/2013/2013-03-20-press-release/index.html

[25] Nikoloulopoulos, A., Joe, H., & Li, H. (2013). Vine copulas with asymmetric tail dependence and applications to financial return data.

[26] Plenum Insurance Linked Capital (February 2014). Annual Review 2013: CAT Bond Market and Insurance Events.

[27] PR Newswire (June 13, 2007). Gamut Re Notes Placed as First Primary Catastrophe Risk CDO. Retrieved March 26, 2015, from http://www.prnewswire.com/news-releases/gamut-re-notes-placed-as-first-primary-catastrophe-risk-cdo-58069052.html

[28] Purnanandam, A. (May 11, 2010). The Originate-to-Distribute Model and the Subprime Mortgage Crisis.

[29] Reuters (April 10, 2014). Manila to conduct new bond swap when conditions right – Purisma. Retrieved March 26, 2015, from http://www.reuters.com/article/2014/04/10/imf-philippines-bonds-idUSL6N0N24G620140410?irpc=932
[30] Risk Management Solutions, Inc. Miu Platform. Retrieved March 26, 2015, from http://www.sifma.org/uploadedfiles/events/2014/insurance_and_risk_linked_securities_conference/riskmanagementsolutions.pdf
[31] SIFMA: Insurance & Risk-Linked Securities Conference (March 19, 2008). Managed Catastrophe Risk Obligations & Dedicated Cat Funds. Retrieved on March 26, 2015, from http://events.sifma.org/uploadedFiles/Events/2008/insurance-risk-linked08/Managed%20Catastrophe%20Risk%20Obligations%20Workshop.Panel.pdf
[32] Standard &Poors. Retrieved March 26, 2015, from http://www.standardandpoors.com/prot/ratings/articles/en/us/?articleType=HTML&assetID=1245361956898 (requires login).
[33] Swiss Re (January 2014). Insurance Linked Securities Market Update.
[34] Trading Risk (July/August 2011). UBS to offer debt to lever ILS funds. Issue Number 36. Retrieved March 26, 2015, from http://cms.insuranceinsider.com/assets/_files/documents/jul_11/ii__1311157133_TR_July_August_LR.pdf
[35] Venter, G. Tails of Copulas.
[36] Willis Re Capital Markets (January 2014). ILS Market Update – Q4: Rounding Off a Year of Plenty.
[37] Woo, G. (February 6-7 2004). A Catastrophe Bond Niche: Multiple Event Risk. NBER Insurance Workshop, Cambridge, Massachusetts.

Abbreviations and notations

Collect here in alphabetical order all abbreviations and notations used in the paper

CDO, collateralized debt obligation
CRO, collateralized risk obligation
EP, exceedance probability
GSAMP, Goldman Sachs GSAMP Trust 2006-NC2
ILS, insurance-linked securities
ILW, industry loss warranty
S&P, Standard and Poor's
SPV, special purpose vehicle

Biography of the Author

Aaron C. Koch, FCAS, MAAA is a Consulting Actuary with Milliman, Inc. Aaron works with a diverse set of clients in the property and casualty industry, including multiline insurers, reinsurers, captives, risk-retention groups, and municipalities. His experience includes reserving and ratemaking for commercial lines of insurance such as property, general liability, workers' compensation, products liability, and medical malpractice.

Aaron develops innovative solutions for the rapidly growing catastrophe risk and alternative capital markets. He works with leaders in this field—including specialist hedge funds and reinsurers—to provide independent analysis and establish best-practice standards for asset valuation, post-event loss reserving, and operational reviews. In addition, Aaron is a frequent writer and speaker on alternative financing structures for catastrophe risk.

Aaron may be contacted via email (aaron.koch@milliman.com) and phone (781-213-6272).