

# Pricing Industry Loss Warranties in a Lévy-Fraily Framework

Simone Beer\*    Alexander Braun†    Andrin Marugg‡

September 2019<sup>§</sup>

## Abstract

We propose a novel risk-neutral pricing approach for industry loss warranties. In doing so, we explicitly take into account the statistical dependence of the losses on individual policies in the underlying insurance portfolio, caused by the occurrence of a natural catastrophe. Inspired by recent advances in the structured credit literature, we model joint claim events in a Lévy-Fraily framework with a stochastic time change. Event time is driven by rare and large jumps of a compound Poisson subordinator and thus elapses more quickly when a natural catastrophe has struck, leading to a clustering of losses. We estimate the model on historical ILW quotes and obtain encouraging fit statistics.

**Key words:** Natural Catastrophe Risk · Industry Loss Warranties · Lévy-Fraily Model · Stochastic Time Change

**JEL Classification:** G13 · G22

---

\*Simone Beer (simone.beer@unisg.ch) has conducted her PhD in Finance at the University of St. Gallen.

†Alexander Braun (alexander.braun@unisg.ch) is with the Institute of Insurance Economics, University of St. Gallen, Tannenstrasse 19, CH-9000 St. Gallen.

‡Andrin Marugg has completed the Master's Program in Banking and Finance at the University of St. Gallen.

§We are grateful to Jan-Frederik Mai as well as two anonymous referees for reviewing our paper and providing helpful comments and suggestions. This research has been funded by the Swiss National Science Foundation under grant number 100018\_179003.

# 1. Introduction

After the landfall of Hurricane Andrew in 1992, the Chicago Board of Trade (CBOT) introduced catastrophe futures and options based on an industry loss index provided by Property Claims Services (see, e.g., Kielholz and Durrer, 1997). Despite the fact that trading in these new risk management instruments ceased a few years later, the era of catastrophe risk securitization had dawned. In 1994, the year of the disastrous Northridge earthquake, Hannover Re launched the first catastrophe bond (cat bond), a fully collateralized solution that allowed the trading of natural disaster risk in a classical fixed income format (PartnersRe, 2015). Shortly after, in 1995, Nationwide issued USD 400 mn of so-called “Act of God Bonds” in a deal arranged by J.P. Morgan and Salomon Brothers (see, e.g., Jaffee and Russell, 1997). Since then, the market for insurance-linked securities (ILS) has been growing markedly (see, e.g., Braun and Weber, 2017). According to Aon Benfield, 2019, outstanding risk capital reached a new record high of USD 97 bn in 2018 and a recent report of auditing and consulting firm Ernst & Young even predicts the ILS market volume to amount to USD 224 bn or 28% of the global reinsurance capital by 2021.

The extant literature has dealt with various aspects of ILS such as sponsor and investor demand, market development, and relation to classical reinsurance (see, e.g., Bantwal and Kunreuther, 2000; Cummins and Weiss, 2009a; Braun et al., 2013; Braun and Weber, 2017). One of the key questions in the context of ILS, however, remains their pricing.<sup>1</sup> Although the number of articles on this issue has notably increased over the last years, a commonly accepted convention is missing to date. Instead, the literature forked into various strands, advocating actuarial, econometric, utility-based, and risk-neutral approaches (see, e.g., Braun, 2016). We aim to contribute to the resolution of this disarray by advancing the debate in favor of the latter category. The general applicability of arbitrage methods in the context of catastrophe-linked instruments has already been evaluated by Balbás et al. (1999) and, more recently, Gatzert et al. (2017) as well as Beer and Braun (2019). Notable attempts to derive risk-neutral pricing formulas were, amongst others, undertaken by Bakshi and Madan (2002) as well as Jarrow (2010). In both cases, the underlying insurance losses are modeled in an aggregate manner. The former studied the valuation of Asian catastrophe loss options based on a mean-reverting Markov process. The latter introduced a reduced-form model for cat bonds, which takes advantage of their similarities with credit-linked notes.<sup>2</sup>

---

<sup>1</sup>While our focus is on nonlife ILS, this also holds for traded life insurance risks (see, e.g., Braun et al., 2012, 2016).

<sup>2</sup>Hence, a cat bond can be essentially decomposed into a default-free bond and a catastrophe swap.

In contrast to most earlier work, we depart from the common practice of aggregate loss modeling. That is, we explicitly consider the stochastic dependence of individual claims in the underlying insurance portfolio. Those are then cumulated to determine whether a payment on the ILS instrument is triggered. Although, at first glance, cat bonds are a natural subject matter for our study, their lack of standardization makes them less suitable for calibration purposes. More specifically, the wide variety of terms and conditions leads to a high-dimensional setting with only few observations for each case (Beer and Braun, 2019). In addition to this data scarcity problem, cat bond issuances with more than one tranche are generally quite rare (see, e.g., Guy Carpenter, 2016). This implies that an aggregate loss process is regularly sufficient for pricing purposes. Therefore, we focus on industry loss warranties (ILWs), which distinguish themselves through complete standardization. Moreover, they cover multiple layers of one and the same portfolio loss distribution. Pricing those with a small set of common parameters logically demands a model that accounts for the stochastic dependence of the underlying risks. With respect to pricing, ILWs have received little scholarly attention to date. Two exceptions are Gatzert et al. (2011) and Gatzert and Schmeiser (2012). The former introduce a simple option-theory framework based on geometric Brownian motions and the latter contrast actuarial and financial valuation. Moreover, Braun (2011) acknowledges the structural similarities of catastrophe swaps and ILWs and proposes a contingent claims pricing approach that he fits to historical data.

Inspired by recent progress in the area of structured credit modeling, we adapt and extend the Lévy-Frailty framework of Mai and Scherer (2009a,b) to ILWs. This novel approach allows us to capture the excess clustering of insurance claims after a natural disaster through a stochastic time change, driven by a compound Poisson subordinator with rare and large jumps. In addition to the theoretical advancement of the ILW literature, we contribute a rare empirical analysis to this field of research. More specifically, we fit our model to a comprehensive time-series-cross-section data set of historical quotes for US wind and US earthquake contracts, underlining its suitability with excellent results. Although ILWs are an ideal application, the framework is not sensitive to contract specification and can thus be flexibly deployed for the market-consistent valuation of catastrophe risk in general. In other words, by matching all observable quotes for standardized contracts with a small set of parameters, we are able to estimate unobservable prices of more complex instruments in accordance with their natural hedges.

The rest of this paper is structured as follows. In Section 2, we briefly recap the main characteristics of ILW contracts and highlight that the stochastic dependence of the underlying insurance claims is relevant for valuation. Section 3 is centered on the model development. It begins with a brief review of the reduced-form model by Jarrow (2010), which will serve as a benchmark, before extensively laying out the Lévy-Frilty framework of Mai and Scherer (2009a,b) and extending it to dependent claim events. Furthermore, in Section 4, we calibrate the model based on historical ILW prices, assess its fit and interpret the parameter estimates. Finally, in Section 5, we draw our conclusion.

## 2. An ILW Primer

### 2.1 Contract Design

The market for ILWs emerged in the late 1980s. In its classical form, an ILW is a simple and standardized double-trigger reinsurance contract. At inception, the protection buyer pays a premium in return for which he receives a lump sum payment (limit) from the protection seller, if a specified event occurs. The event is defined as a natural disaster that drives an index of insurance industry losses beyond an attachment point (warranty) and, at the same time, causes losses in the insurance portfolio of the protection buyer that exceed a certain threshold (retention). These conditions for the payoff are called index trigger and indemnity trigger, respectively.<sup>3</sup> The latter is set fairly low so that it will almost certainly be hit if industry losses are high enough for a payout. ILWs reference a specific territory (e.g., Florida, California) as well as peril (e.g., hurricanes, earthquakes), and typically exhibit a contract term of 12 months. Moreover, they can be structured as aggregate or occurrence contracts. The former allow for an aggregation of losses caused by all covered natural disasters that strike in the reference territory during the contract term and thus resemble a *stop-loss reinsurance contract*. The latter, in contrast, are akin to *excess of loss per event (CatXL) reinsurance contracts* that can only be triggered by losses that stem from a single catastrophe. From a protection buyer's point of view, the major drawback of ILWs is their basis risk, arising from the fact that insurance losses incurred by the whole industry are not perfectly correlated with those of the protection buyer (see, e.g., Zeng, 2000).

---

<sup>3</sup>Due to the indemnity trigger, classical ILWs qualify for reinsurance accounting (see, e.g., Cummins and Weiss, 2009b).

Apart from the classical double-trigger ILW, there is a second form that exclusively exhibits an industry loss trigger. The U.S. firm Property Claims Services (PCS) provides the dominant index suite for such transactions. Since an own insurance portfolio is not required to purchase this product, it is essentially a derivative instead of a reinsurance contract. Due to their structural similarities with *credit default swaps* (CDS), practitioners dub ILWs without an indemnity trigger *catastrophe swaps* (see, e.g., Braun, 2011). The latter are mainly employed by alternative investment funds for the tactical management of their ILS portfolios (see, e.g., Braun et al., 2019a). For a more comprehensive review of ILWs, the reader is referred to McDonnell (2002) as well as Gatzert and Schmeiser (2012).

## 2.2 The Relevance of Stochastic Dependence in ILW Pricing

Henceforth we focus on ILWs in the derivatives form (catastrophe swaps), since the absence of an indemnity trigger facilitates risk-neutral valuation. Multiple layers of ILWs with different attachment points stacked on top of each other are essentially quite similar to the tranche structure of a collateralized debt obligation (CDO) or an asset-backed security (ABS). In the same way as default losses in a pool of mortgages, loans or bonds are absorbed by the tranches in the capital structure of a CDO, claims in an insurance portfolio may be redistributed to reinsurance or ILW layers according to a predefined hierarchy. Of particular interest in this regard is the pool of insurance policies that underlies the industry loss index referenced by the ILW contract. Since it reflects the aggregated catastrophe event losses of the entire insurance industry, the index is essentially an enormously large portfolio, consisting of all policies that have been underwritten in the covered territory. As a result, the shape of the relevant loss distribution and, in turn, the value of each ILW layer, is driven by the stochastic dependence of the individual risks in the index portfolio. For the sake of argument, we will center the following explanations on correlation. Note, however, that our model framework does not rely on the assumption of linear dependencies.<sup>4</sup>

Under normal circumstances, property insurance risks are virtually uncorrelated. To see this, consider an insurance policy on a house in the south and in the north of Florida. Usually, one has no reason to expect a loss on both at the same time. In other words, the risk of fire, burglaries, or other claim events

---

<sup>4</sup>Mai (2010) derives the implied dependence structure of the Lévy-Frilty default model and shows that it belongs to the Marshall-Olkin family.

is idiosyncratic to each property. Consequently, large policy pools benefit from supreme diversification effects. The latter form the core of the insurance business model. Expected losses are priced into the premiums and the variation around them tends to be small enough to be covered by risk loadings and equity capital. This general rule, however, breaks down in the presence of loss accumulation scenarios. Those are caused by large-scale events such as extreme natural disasters (e.g., Hurricane Irma in 2017), which lead to correlated losses on originally independent insurance policies (see Braun et al., 2019b). In these scenarios, the aggregate (portfolio-level) loss distribution exhibits a substantial positive skewness, implying that diversification largely breaks down (see, e.g., Jaffee and Russell, 1997). Most insurers need to cede such tail risks in order to avoid existence-threatening capital events.<sup>5</sup>

### 3. Pricing ILWs

#### 3.1 Reduced-Form Model

Consider a probability space  $(\Omega, \mathcal{F}, \mathbb{Q})$ . We explicitly specify the model under the risk-neutral measure  $\mathbb{Q}$ . Much like a CDS, an ILW (catastrophe swap) can be priced by equating the present values of the protection and the premium leg. Since the premium, denoted by  $p^{\text{ILW}}$ , is typically paid upfront and no further payments are made by the protection buyer to the protection seller, the problem reduces to valuing the *protection leg*. The present value (PV) of the latter equals the expected discounted stream of stochastic future payoffs of the contract (from the protection seller to the protection buyer). For ease of calculation, we fix a schedule of discrete dates  $\mathcal{M} = [t_1 \ t_2 \ t_3 \ \dots \ t_T]$  on which a trigger event may occur.  $T$  stands for the final maturity of the contract. Provided the differences  $t_i - t_{i-1}$  are kept reasonably small, the PV of the *protection leg* can be approximated as follows (see O’Kane and Turnbull, 2003):

$$p^{\text{ILW}} \approx L \sum_{i=1}^T DF(t_i) (TP(t_i) - TP(t_{i-1})), \quad t_i \geq 0; \quad (1)$$

where  $L$  is the specified limit (the lump-sum payoff) of the ILW,  $DF(t_i)$  represents the discount factor that matches time  $t_i$  and  $TP(t_i)$  denotes the risk-neutral probability that the contract is triggered at or

---

<sup>5</sup>Correlated risks in the insurer’s investment and underwriting portfolios lead to higher ruin probabilities and therefore substantial markups in regulatory capital charges. For an illustration of this effect on the asset side of the balance sheet see, e.g., Braun et al. (2015, 2017, 2018).

before time  $t_i$ .<sup>6</sup> Accordingly,  $TP(t_i) - TP(t_{i-1})$ , also denoted  $TP(t_{i-1}, t_i)$ , equals the probability of a trigger event occurring between times  $t_{i-1}$  and  $t_i$ . The following parameters are required to calculate  $p^{\text{ILW}}$ : (i) the set of risk-neutral trigger probabilities, (ii) the deterministic limit of the contract, and (iii) the zero-coupon yield curve to discount cash flows. Items (ii) and (iii) are readily observable. Hence, the only non-trivial quantities are the risk-neutral trigger probabilities. The issue of modeling them has, amongst others, been addressed by Embrechts and Meister (1997), Geman and Yor (1997), or Bakshi and Madan (2002) in various catastrophe derivative contexts.

More recently, Jarrow (2010) suggested to price catastrophe bonds by means of a standard reduced-form pricing model from the credit risk literature (see, e.g., Jarrow and Turnbull, 1995). Below we adapt his approach to ILWs and use it as a benchmark for the Lévy-Frailty framework advocated in this paper. Assume that trigger events are driven by a nonhomogeneous Poisson process  $N_t$  with intensity  $\lambda_{t_{i-1}, t_i}$ , defined as the expected number of events between times  $t_{i-1}$  and  $t_i$  under the risk-neutral measure  $\mathbb{Q}$ :

$$\lambda_{t_{i-1}, t_i} = \mathbb{E}_{t_{i-1}}^{\mathbb{Q}} (N_{t_i} - N_{t_{i-1}}) = \int_{t_{i-1}}^{t_i} \lambda(u) du.$$

Hence, the contract's survival probability  $SP(t_{i-1}, t_i) = 1 - TP(t_{i-1}, t_i)$  between times  $t_{i-1}$  and  $t_i$  ( $t_i > t_{i-1}$ ) equals

$$SP(t_{i-1}, t_i) = \Pr(N_{t_i} - N_{t_{i-1}} = 0) = \frac{e^{-\lambda_{t_{i-1}, t_i}} (\lambda_{t_{i-1}, t_i})^0}{0!} = e^{-\lambda_{t_{i-1}, t_i}},$$

and the risk-neutral probability for a trigger event between  $t_{i-1}$  and  $t_i$  can be described as follows:

$$TP(t_{i-1}, t_i) = \Pr(N_{t_i} - N_{t_{i-1}} > 0) = 1 - SP(t_{i-1}, t_i) = 1 - e^{-\lambda_{t_{i-1}, t_i}}. \quad (2)$$

Below, we deviate from such an aggregate trigger process by explicitly considering the individual items in the insurance portfolio on which the contract is written. Although our work focuses on ILWs, the idea is easily extendable to other types of catastrophe derivatives. To avoid any confusion regarding terminology,

---

<sup>6</sup>As discussed in the previous section, we will exclusively focus on ILWs in derivatives form. Recall that, in this case, a trigger event is defined as the occurrence of a catastrophe in the covered territory that causes insurance industry losses in excess of the contract's attachment point. On the flip side, survival of the contract implies that losses remain below the attachment point.

we will refer to the occurrence of a natural disaster (e.g. hurricane, earthquake) as a *catastrophe event* and to a loss on a specific insurance policy as a *claim event*.

### 3.2 The Marshall-Olkin and the Lévy-Frilty Framework

The Marshall-Olkin (MO) copula naturally extends the ideas of single-name reduced-form credit risk models to the multivariate case. The first MO approach in financial economics was developed by Lindskog and McNeil (2003), who distinguishes between shocks that are idiosyncratic to the firm (e.g. fraud) or systematic across the economy (e.g. recessions). His method can be employed to price both basket credit derivatives and portfolios of insurance lines. However, an important drawback of the model is its numerical instability in high-dimensional settings (see, e.g., Mai, 2010). This is a major weakness in the context of structured credit index tranches and even more so in the context of very large insurance pools.

Therefore, we resort to a modern extension of the classical MO model: the Lévy-Frilty default framework introduced by Mai and Scherer (2009a,b). Relying on a stochastic time change in combination with the large homogeneous portfolio (LHP) assumptions, the latter overcomes the aforementioned obstacles and exhibits a high degree of mathematical tractability. In analogy to joint credit defaults during economic downturns, systematic insurance risks such as large-scale natural disasters lead to dependent claim events. A Lévy subordinator that serves as a stochastic clock can capture this phenomenon, since it allows time to elapse more quickly in the aftermath of a catastrophe, causing losses to cluster. As will be shown below, further advantages of the Lévy-Frilty framework are a clean separation of the marginal distributions and the dependence structure as well as an intuitive economic interpretation of the parameters for the time change. Furthermore, the implied copula that drives joint survival probabilities can be computed explicitly.<sup>7</sup> Finally, the LHP assumptions are highly reasonable for ILWs. The pool of policies underlying a loss index as calculated by PCS is extremely large, since it consists of all portfolios held by the insurance industry as a whole. Likewise, homeowners insurance contracts are quite homogeneous. According to data from the National Association of Insurance Commissioners (NAIC), 78% of the customers in the U.S. opt for an open-peril HO-3 policy, which covers any cause of damage apart from a few explicit exclusions. In addition, the insured sums in residential property portfolios are usually distributed quite evenly.

---

<sup>7</sup>Mai and Scherer (2009a) show that in the bivariate case, the implied copula is of Cuadras-Augé type.



### 3.3 Stochastic Time Change and Subordinators

The concept of a stochastic time change was first proposed by Clark (1973), who noticed that cotton futures prices exhibit a higher volatility on days with large transaction volumes. However, when measured with respect to trading activity, prices are normally distributed. This finding led other authors to suggest that changing the time measure might be a suitable approach for the modeling of asset prices (see, e.g., Ané and Geman, 2000; Geman, 2002; Carr and Wu, 2004). Put differently, it is possible to define a non-equidistant time grid such that asset returns have the same distribution in each interval (see, e.g., Borovkova and Schmeck, 2017). The modified time measure is called *business time* and can be interpreted as the arrival of new information (see, e.g., Tankov, 2003).

To be suitable for the modeling of business time, stochastic processes need to have right-continuous sample paths with left limits, and exhibit both stationary and independent increments. Lévy subordinators fulfill all of these requirements (see, e.g., Tankov, 2003). Joshi and Stacey (2006) were the first to apply them for the generation of dependent credit events. During normal times, there is only a small number of defaults. However, when the economy enters a downturn, more firms become financially distressed. Consequently, a lot of business time should elapse during a recession. Chang et al. (1996, 2008, 2010) introduce a similar approach for catastrophe futures and options. The case of a whole spectrum of ILW layers, however, has not been considered yet. This is surprising, given that ILW prices are evidently driven by the stochastic dependence of the risks in the underlying index portfolio (refer to the second section). Below, we extend the work of Mai and Scherer (2009a) to address this issue. Consistent with their model, claim dependence among individual policies will be introduced through a common stochastic time change based on subordinators. We generally follow their notation and highlight our own model alterations explicitly.

### 3.4 Modeling Isolated Claim Events

We begin with the specification of the marginal claim event distributions. Consider a pool of  $n$  insurance policies with claim times  $\tau_i$  ( $i = 1, 2, \dots, n$ ). Under normal circumstances, claim events are stochastically independent. Thus, it is reasonable to assume that they are driven by a homogeneous Poisson process  $N_i(t)$  with hazard rate function  $h_i(t) = \lambda_i t$  and constant and individual hazard rate  $\lambda_i$  idiosyncratic to

policy  $i$  ( $\lambda_i > 0$ ). A claim event for policy  $i$  coincides with the first jump of the Poisson process:

$$\tau_i = \inf\{t \geq 0 : N_i(t) > 0\}, \quad \forall i = 1, 2, \dots, n; \forall t \geq 0. \quad (3)$$

The probability of a loss on policy  $i$  at or before time  $t$ ,  $\Pr(\tau_i \leq t) = POL_i(t)$ , is defined as:

$$\begin{aligned} POL_i(t) &= 1 - PNL_i(t) \\ &= 1 - e^{-h_i(t)} \\ &= 1 - e^{-\lambda_i t}, \quad \forall i = 1, 2, \dots, n; \forall t \geq 0, \end{aligned} \quad (4)$$

where  $\Pr(\tau_i > t) = PNL_i(t)$  equals the probability that policy  $i$  suffers no loss at or before time  $t$ . Although the approach is well-suited to model the occurrence of individual events, it cannot capture the joint behavior of multiple policies in a pool. For this purpose, we need to resort to the MO framework for dependent events.

### 3.5 Modeling Simultaneous Dependent Claim Events

The MO framework extends the conventional reduced-form approach by allowing the shock to have multiple reasons. Consider a setup with two shock processes of Poisson type. In this case, the claim time  $\tau_i$  for policy  $i$  on the time scale  $t$  is defined as

$$\tau_i = \inf\{t \geq 0 : N_i(t) + N(t) > 0\}, \quad \forall i = 1, 2, \dots, n; \forall t \geq 0, \quad (5)$$

where  $N_i(t)$  is a Poisson process with intensity  $\lambda_i$  idiosyncratic to policy  $i$  and  $N(t)$  represents the common shock process with intensity  $\lambda$  ( $\lambda_i > 0; \lambda > 0$ ). Consequently, policy  $i$  experiences a claim event if either one of the two shocks occurs for the first time. Owing to the superposition property of Poisson processes, the intensity at which policy  $i$  suffers a loss equals the sum of the default intensities of the individual processes.<sup>8</sup> The probability of no loss on policy  $i$  at or before time  $t$ ,  $\Pr(\tau_i > t) = PNL_i(t)$ ,

---

<sup>8</sup>A superposition of independent Poisson processes is again a Poisson process. To put it formally, if  $N_1(t)_{t \geq 0}$  and  $N_2(t)_{t \geq 0}$  are independent Poisson processes with intensities  $\lambda_1$  and  $\lambda_2$ , then  $(N_1(t) + N_2(t))_{t \geq 0}$  is a Poisson process with intensity  $\lambda_1 + \lambda_2$  (see, e.g., Tankov, 2003).

can thus be defined as follows:

$$PNL_i(t) = e^{-(\lambda_i + \lambda)t}, \quad \forall i = 1, 2, \dots, n; \forall t \geq 0. \quad (6)$$

Further systematic risk factors can be added as one sees fit. While the choice of copula is often flexible, this model requires the marginal distributions to be linked into a joint survival function by means of the MO copula known from reliability theory (see Marshall and Olkin, 1967). If there are two items with individual survival functions as specified in Equation (6), a unique function  $C^\tau: [0, 1]^2 \rightarrow [0, 1]$  exists through which the joint survival probabilities can be represented as follows:<sup>9</sup>

$$\Pr(\tau_1 > t, \tau_2 > t) = C^\tau(PNL_1(t), PNL_2(t)). \quad (7)$$

Now let  $\theta_i$  denote the ratio of the common intensity  $\lambda$  to the total intensity  $(\lambda_i + \lambda)$  for policy  $i$ , i.e.

$$\theta_i = \frac{\lambda}{\lambda_i + \lambda}. \quad (8)$$

In this case, the MO copula can be represented as follows:

$$C^\tau(PNL_1(t), PNL_2(t)) = \min(PNL_1(t) \cdot PNL_2(t)^{1-\theta_1}, PNL_2(t) \cdot PNL_1(t)^{1-\theta_2}). \quad (9)$$

This bivariate model can be easily extended to the multivariate case.<sup>10</sup>

### 3.6 Modeling Non-Simultaneous Dependent Claim Events

Due to the common shock in the MO model, claim events caused by the systematic risk factor(s) always occur simultaneously. We relax this assumption by switching from conventional time  $t$  to *event time*  $S_t$ .<sup>11</sup> The latter will be governed by an independent Lévy subordinator  $S = \{S_t\}_{t \geq 0}$  (almost surely non-

---

<sup>9</sup>The theoretical foundation is provided by Sklar's theorem.

<sup>10</sup>Giesecke (2003), for example, draws on the MO framework to price basket default swaps.

<sup>11</sup>Owing to the close link between credit events and subsequent business disruptions, Mai and Scherer (2009a,b) term the subordinator *business time*. Since the latter may not be considered as an adequate expression for modeling catastrophe events, we decide to call it *event time* instead.

decreasing) with characteristic triplet  $(\mu, 0, \nu)$ .<sup>12</sup> Hence, it is completely determined by its corresponding Laplace transform:

$$\mathbb{E}[e^{-\lambda_i S_t}] = e^{-\psi(\lambda_i)t}, \quad \forall i = 1, 2, \dots, n; \forall t \geq 0, \quad (10)$$

where  $\psi$  represents the Laplace exponent and is defined as follows:

$$\psi(\lambda_i) := \mu\lambda_i + \int_0^\infty (1 - e^{-\lambda_i t}) \nu(dt), \quad \forall i = 1, 2, \dots, n. \quad (11)$$

We focus on a single risk source as in Equation (4). After the time change, the intensity  $\lambda_i$  represents the rate at which policy  $i$  experiences a claim event per unit of event time  $S_t$ . Thus, the original claim event probability in Equation (4) changes as follows:

$$POL_i(t) = 1 - \mathbb{E}[e^{-\lambda_i S_t}], \quad \forall i = 1, 2, \dots, n; \forall t \geq 0. \quad (12)$$

Now, claim time  $\tau_i$  of policy  $i$  can be interpreted as the first jump of a homogeneous Poisson process subject to a common event time. As a consequence, we have stochastically dependent events with non-simultaneous occurrences. For a formal definition as well as key distributional properties, please refer to Mai and Scherer (2009a). The latter also demonstrate how the marginal distributions can be kept invariant under the time change to ensure analytical tractability. In our context,  $S$  satisfies the so-called time normalization for the exponential distribution, if the following relationship holds (proof see Mai and Scherer, 2009a):

$$\mathbb{E}[F(S_t)] = F(t), \quad \forall t \geq 0, \quad (13)$$

such that

$$1 - \mathbb{E}[e^{-\lambda_i S_t}] = 1 - e^{-\lambda_i t}, \quad \forall i = 1, 2, \dots, n; \forall t \geq 0. \quad (14)$$

Put differently, the parameter values of the subordinator have to be chosen so as to ensure Equation (14)

---

<sup>12</sup>Lévy processes are fully specified by their characteristic triplet  $(\mu, \sigma, \nu)$ .  $\mu$  and  $\sigma$  denote the drift and volatility of a Wiener process, while the Lévy measure  $\nu$  reflects size and frequency of jumps.  $\nu([x, \infty))$  gives the expected number of jumps per time step that are greater than or equal to  $x$ .

to hold. This applies if  $-\psi(\lambda_i) = -\lambda_i$ .<sup>13</sup> We thus achieve time normalization by constructing the following deterministic relationship between  $\mu$  and  $\nu$ :

$$\mu_i = 1 - \frac{1}{\lambda_i} \int_0^\infty (1 - e^{-\lambda_i t}) \nu(dt), \quad \forall i = 1, 2, \dots, n. \quad (15)$$

In the case of ILWs, only one type of systematic insurance risk (e.g., windstorms or earthquakes) is relevant. All other risk sources are excluded from the instrument by construction (see second section). Nevertheless, depending on their geographic locations, different insured objects in the underlying insurance portfolio can be more or less exposed to relevant natural disasters. Hence,  $\lambda_i$  in Equation (12) stands for the intensity of policy  $i$  with regard to the single source of catastrophic risk that is covered by the ILW contract at hand.

### 3.7 Deriving the Aggregate Insurance Loss Distribution Function

As in the case of structured credit products, the key input factor for the pricing of a whole spectrum of ILW layers is the cumulative distribution function (*cdf*) of the aggregate portfolio loss. Below, we describe how the latter can be derived in the Lévy-Frailty framework. In our insurance portfolio with  $n$  constituents, let  $L_t^n$  be the proportion of policies that suffered losses up to time  $t$ . Then the relative loss process  $L^n = \{L_t^n\}_{t \geq 0}$  can be expressed as follows:

$$L_t^n := \frac{1}{n} \sum_{i=1}^n A_t^i, \quad \forall t \geq 0. \quad (16)$$

The stochastic process  $A^i = \{A_t^i\}_{t \geq 0}$  is defined as:

$$A_t^i := \mathbf{1}_{\tau_i \leq t}, \quad \forall i = 1, 2, \dots, n; \forall t \geq 0. \quad (17)$$

For any  $t$ , each  $A_t^i$  follows a Bernoulli distribution with probability of success  $POL_i(t) = 1 - \mathbb{E}[e^{-\lambda_i S_t}]$ .

---

<sup>13</sup>Proof:

$$\begin{aligned} 1 - \mathbb{E}[e^{-\lambda_i S_t}] &= 1 - e^{-\lambda_i t} \\ \Leftrightarrow 1 - e^{-\psi(\lambda_i)t} &= 1 - e^{-\lambda_i t} \\ \Leftrightarrow -\psi(\lambda_i) &= -\lambda_i. \end{aligned}$$

If the portfolio is homogeneous, the claim times  $\tau_i$  exhibit the same distribution. Consequently, all intensities must be identical and we may write  $\lambda_i = \lambda$ . Mai and Scherer (2009a) show that, under this assumption, the absolute number of impaired policies  $nL_t^n$  can be approximated by means of the Bernoulli mixture model

$$nL_t^n \sim \text{Bin}(n, 1 - e^{-\lambda S_t}), \quad \forall t \geq 0. \quad (18)$$

In addition, as the portfolio size  $n$  tends to infinity,  $L_t^n$  converges to the mixture variable

$$\lim_{n \rightarrow \infty} L_t^n = L_t^\infty = 1 - e^{-\lambda S_t}, \quad \forall t \geq 0. \quad (19)$$

Thus, if the LHP assumptions hold, we exactly know the relative loss  $L_t^\infty$ , given any realization of the stochastic process  $S_t$ . It is thus straightforward to derive the probability density function (*pdf*) of the random percentage of impaired policies in the portfolio,  $f_{L_t^\infty}(x)$ , from the *pdf* of the subordinator,  $f_{S_t}(x)$ :<sup>15</sup>

$$f_{L_t^\infty}(x) = f_{S_t}\left(-\frac{1}{\lambda} \log(1-x)\right) \frac{1}{\lambda(1-x)}, \quad x \in (0, 1); \forall t \geq 0. \quad (20)$$

### 3.8 Introducing Recovery Rates

So far we have implicitly assumed that a claim event will inevitably lead to a full loss of a policy's sum insured, i.e., Equation (16) is a zero-recovery loss process. In reality, however, insurance companies will often face partial losses, for example, due to deductibles or because the value of the covered object has not been entirely destroyed. Hence, the characteristics of the industry index portfolio referenced by an ILW evidently have an influence on the aggregate loss distribution and thus the trigger probabilities. In the LHP setting introduced above, all portfolio constituents are assumed to exhibit the same contract specifications, which leads to identical claim intensities  $\lambda$  and sums insured  $SI$ . We may therefore describe the loss on an individual policy  $l_i$ , conditional on the occurrence of a claim event, based on a deterministic recovery rate  $R$  (in %):

$$l_i = l = (1 - R) \cdot SI \quad (21)$$

---

<sup>14</sup>See Mai (2010): “ $\sup_{t \geq 0} |L_t^n - (1 - e^{-\lambda S_t})|$  tends to zero almost surely.”

<sup>15</sup>Consider two random variables  $X$  and  $Y$  with known densities, where  $Y = G(X)$  and  $G(X)$  is non-decreasing. Then  $f_Y(y) = f_X(G^{-1}(y)) \frac{dG^{-1}(y)}{dy}$ . In our case:  $Y = L_t^\infty$ ,  $X = S_t$  and  $G^{-1}(y) = -\frac{1}{\lambda} \log(1-y)$ .

$R$  is, by definition, independent of the loss process. Yet, it will vary with the combination of peril (windstorm, earthquake etc.) and geography (e.g., United States, Japan etc.) covered by an ILW and can either be determined exogenously or estimated during the calibration process. Mai and Scherer (2009a) demonstrate that, under these circumstances, Equation (19) becomes

$$L_t^{\infty,R} = (1 - R) (1 - e^{-\lambda S_t}), \quad \forall t \geq 0; \quad (22)$$

and the corresponding *pdf* is

$$f_{L_t^{\infty,R}}(x) = f_{S_t} \left( -\frac{1}{\lambda} \log \left( 1 - \frac{x}{1-R} \right) \right) \frac{1}{\lambda(1-R-x)}, \quad x \in (0, 1); \forall t \geq 0. \quad (23)$$

Alternatively, one may implement the model based on Equation (20) and take into account a fixed portfolio-level recovery rate  $R$  by adjusting the integration boundaries for the calculation of the expected tranche loss. For an illustration of this method in the context of structured credit instruments, the interested reader is referred to Kalemanova et al. (2007).

### 3.9 Selecting a Suitable Subordinator

An explicit solution of the *pdf* of the relative loss, as shown in Equations (20) and (23), calls for a specification of  $S_t$ . Here, we must account for the fact that catastrophe events are rare, but cause insurance claims to occur in great numbers in a very short time interval. To reflect this pattern, we require a subordinator with infrequent and large jumps. Infinite activity subordinators and drift terms, as employed by Mai and Scherer (2009a), are not appropriate, as they imply that event time is always moving. In contrast, we opt for a compound Poisson subordinator, which is of the convenient form:

$$S_t = \sum_{j=1}^{M(t)} Y_j, \quad \forall t \geq 0, \quad (24)$$

where  $M(t)$  stands for a Poisson process with intensity  $\beta$ . The  $Y_j$  are i.i.d. jump sizes.<sup>16</sup> To fulfill the requirements of a subordinator, the jump sizes must adhere to a distribution which is exclusively defined on the positive half axis. We will return to this issue below. It can be shown that the Laplace exponent

---

<sup>16</sup>Note that for a fixed jump size  $\bar{Y} = 1$ , the compound Poisson process is equivalent to a conventional Poisson process.

$\psi$  of the compound Poisson process (without a drift) has the following form:

$$\psi(u) = \beta \mathbb{E} [(1 - e^{-uY_1})], \quad u \geq 0. \quad (25)$$

Furthermore, time normalization of the subordinator in line with Equation (10) can be achieved by setting  $\psi(1) = 1$ , which leads to the following value for  $\beta$ :

$$\beta = \frac{1}{1 - \mathbb{E}[e^{-Y_1}]}. \quad (26)$$

In this setup, catastrophe events are infrequent but affect a large number of policies simultaneously. Since the subordinator exhibits discontinuous sample trajectories, we have a slightly modified *pdf* of the relative loss process compared to the original  $L_t^\infty$  shown in Equation (20) (see, e.g., Böttcher, 2010):<sup>17</sup>

$$f_{L_t^\infty}(x) = \sum_{k=0}^{\infty} \Pr(M(t) = k) f_k \left( -\frac{1}{\lambda} \log(1-x) \right) \frac{1}{\lambda(1-x)}, \quad x \in (0, 1); \forall t \geq 0. \quad (27)$$

where  $f_k$  is the *pdf* of the  $k^{th}$  convolution of the subordinator's jump size distribution.<sup>18</sup> As we do not need to calculate expected losses to price ILWs, we directly use the *cdf* to speed up computation. It is of the form

$$F_{L_t^\infty}(x) = \sum_{k=0}^{\infty} \Pr(M(t) = k) F_k \left( -\frac{1}{\lambda} \log(1-x) \right), \quad x \in (0, 1); \forall t \geq 0. \quad (28)$$

$F_k$  is the *cdf* of the  $k^{th}$  convolution of the subordinator's jump size distribution. This means that a tractable model is achieved if the  $k^{th}$  convolution is explicitly known (see Mai and Scherer, 2009a). To ensure this, we will discuss adequate choices for the jump size distribution in the next section.

---

<sup>17</sup>This is not necessary for infinite activity subordinators.

<sup>18</sup>The convolution of the individual probability distributions of two or more independent random variables delivers the distribution of the sum of these variables.



Equation (28) specifies a *cdf* for the loss fraction  $x$  given by the attachment point of the ILW in relation to the pool size:

$$x = \frac{\text{attachment point}}{\text{pool size}}. \quad (29)$$

To additionally account for the possibility that not the entire value of an insured item may be destroyed in case a catastrophe occurs, we divide the trigger threshold by the *adjusted pool size*:

$$x = \frac{\text{attachment point}}{\text{adjusted pool size}}, \quad (30)$$

which incorporates the recovery rate  $R$  defined in Equation (21) as follows:

$$\text{adjusted pool size} = (1 - R) \cdot \text{pool size}. \quad (31)$$

Alternatively, we could have also directly modified the portfolio loss *cdf* in Equation (28) as illustrated by Equation (23) for the *pdf* case. We explain our choice to opt for an adjustment of the pool size in more detail, when we discuss the estimates for the parameters.

### 3.10 Selecting the Jump Size Distribution

For the jump size we require a distribution that is exclusively defined on the positive half axis and has a known density under convolution. These two requirements considerably reduce the range of suitable candidates to the Cauchy, chi-squared, exponential, and gamma distribution.<sup>19</sup> The first one needs to be excluded as well, since it has an undefined mean and variance. Consequently, we further explore the  $k$ -th convolutions of the gamma, chi-squared and exponential distribution:<sup>20</sup>

- $\sum_{i=1}^k \text{Gamma}(\alpha, \theta) \sim \text{Gamma}(k\alpha, \theta)$ , with shape  $\alpha > 0$  and scale  $\theta > 0$ ,
- $\sum_{i=1}^k \chi^2(r_i) \sim \chi^2(\sum_{i=1}^k r_i)$ , with degrees of freedom (mean)  $r = \sum_{i=1}^k r_i$
- $\sum_{i=1}^k \exp(\eta) \sim \text{Gamma}(k, \eta)$ , with rate  $\eta > 0$ .<sup>21</sup>

---

<sup>19</sup>The range gets even smaller when considering a continuous-time setting.

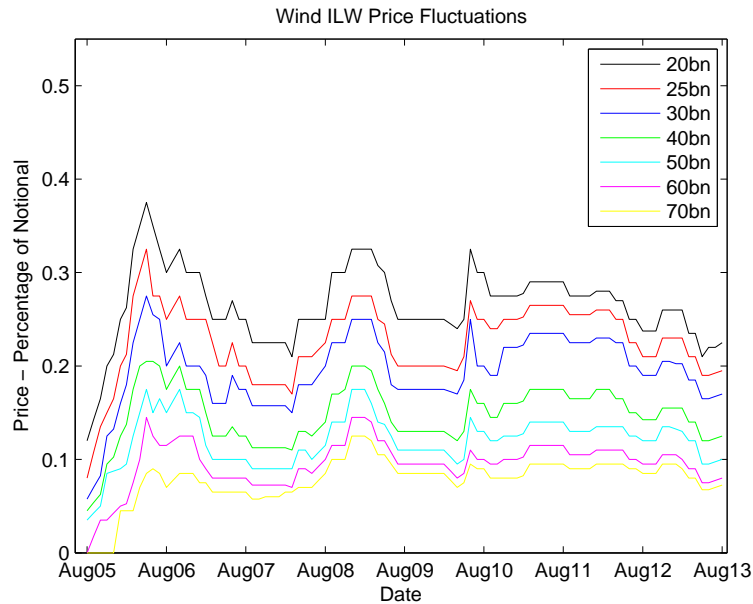
<sup>20</sup>Note that, for these relationships to hold, stochastic independence of the summands is required.

<sup>21</sup>We know that, for the exponential distribution, Equation (26) reduces to  $\beta = 1 + 1/\eta$  (see Mai and Scherer, 2009a).

## 4. Empirical Analysis

### 4.1 Dataset

Our dataset consists of end-of-month ILW prices for US wind and US earthquake contracts in the time period between August 2005 and August 2013. The quotes were obtained from BMS and reflect the premium  $p^{\text{ILW}}$  at the initiation of a new contract with a fixed maturity of 12 months. For both perils, we have seven attachment points (warranties), defined as  $W = \{20, 25, 30, 40, 50, 60, 70\}$  and  $E = \{10, 15, 20, 25, 30, 40, 50\}$  for wind and earthquake contracts, respectively (quoted in USD bn).<sup>22</sup> All prices are percentages of the selected notional. For example, on 31 August 2005, a US wind ILW with a USD 20 bn warranty was quoted at 12%. This means that it cost an up-front payment of USD 12 mn to enter a contract with a notional (payoff) of USD 100 mn as the protection buyer.



**Figure 1: Time series of monthly ILW prices for US wind contracts (08/2005–08/2013)**

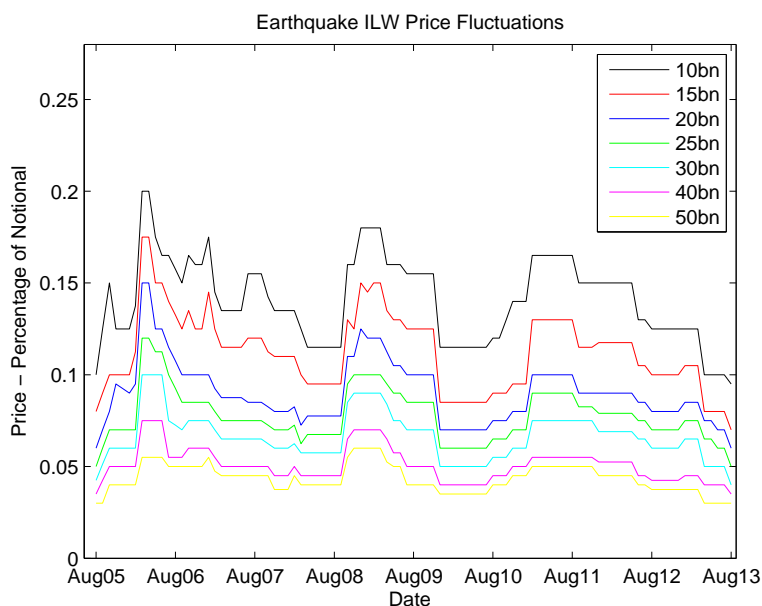
Figure 1 illustrates the development of monthly prices for US wind ILWs between August 2005 and August 2013. The time period (in months) is depicted on the x-axis and the price as a percentage of the notional on the y-axis. The quotes reflect the premium  $p^{\text{ILW}}$  at the initiation of a new contract for a fixed maturity of 12 months. The cross section of attachment points (warranties) is defined as  $W = \{20, 25, 30, 40, 50, 60, 70\}$  (quoted in USD bn).

Figure 1 shows the price movements of US wind ILWs. In August 2005, reinsurance was quite cheap, with the highest two warranties (USD 60 and USD 70 bn) not even quoted and the lowest warranty

---

<sup>22</sup>Recall from the second section that the warranty equals the total industry loss that needs to be exceeded for a payoff.

being just at 12%. Following Hurricane Katrina, prices rose steadily until they peaked in the summer of 2006. This was a classical hard market with high premiums and low reinsurance capacity (see, e.g., Cummins et al., 2006). After a brief recovery phase, during which conditions softened again, a sharp rise was observed around the time of the collapse of Lehman Brothers and the occurrence of Hurricane Ike. The next increase could be observed in the summer of 2010. During that hurricane season, four category 4 storms were observed: Danielle, Earl, Igor and Julia (NOAA, 2011). From then on until August 2013, premiums slowly decreased. The time series of historical prices of US earthquake ILWs in Figure 2 exhibit very similar movements. This suggests that price fluctuations may also be explained by structural effects such as the reinsurance cycle and not merely by manifestations of the reference peril.



**Figure 2: Time series of monthly ILW prices for US earthquake contracts (08/2005–08/2013)**

Figure 2 illustrates the development of monthly prices for US earthquake ILWs between August 2005 and August 2013. The time period (in months) is depicted on the x-axis and the price as a percentage of the notional on the y-axis. The quotes reflect the premium  $p^{\text{ILW}}$  at the initiation of a new contract for a fixed maturity of 12 months. The cross section of attachment points (warranties) is defined as  $E = \{10, 15, 20, 25, 30, 40, 50\}$  (quoted in USD bn).

Apart from the historical ILW price data, we collect end-of-month US zero-coupon yield curves (treasury spot rates) for the time period between August 2005 and August 2013 from Bloomberg. Fitting a Nelson-Siegel-Svensson model allows us to derive monthly discount factors for Equation (1).

## 4.2 Calibration to ILW Prices

### Reduced-Form Model

In line with our ILW and interest rate data, we set the trigger-time vector  $\mathcal{M}$  (see third section) to monthly intervals  $(t_1, t_2, t_3, \dots, t_{12})$ , implying that the time step length  $\Delta t_i = t_i - t_{i-1}$  equals  $\frac{1}{12}$ . Moreover, we switch to a homogeneous Poisson setup by assuming a fixed  $\lambda$  in p.a. terms, which is broken down to match the loss process as follows:  $\lambda_{t_{i-1}, t_i} = \lambda \Delta t_i$  ( $\forall t_i$ ). For each monthly cross-section of wind and earthquake ILW quotes, we then determine the implied intensity rate that minimizes the sum of squared price errors over all seven layers:

$$\operatorname{argmin}_{\hat{\lambda}} \sum_{i=1}^7 \left( \hat{p}_{i,t}^{\text{ILW}}(\hat{\lambda}) - p_{i,t}^{\text{ILW}} \right)^2, \quad (32)$$

where  $\hat{p}_{i,t}^{\text{ILW}}(\hat{\lambda})$  is the price generated by Equation (1) and  $p_{i,t}^{\text{ILW}}$  is the observed market spread. The optimization problem in (32) is solved by means of the Levenberg-Marquardt algorithm. Thus, we obtain one implied intensity estimate  $\hat{\lambda}$  (p.a.) for each monthly cross-section of ILW spreads in our sample.<sup>23</sup>

### Lévy-Fraily Model

Fitting Equation (28) to market data implies the simultaneous estimation of three unknown parameters for the specifications with exponentially and chi-squared distributed jump sizes, and of four parameters for the model with gamma-distributed jump sizes. We proceed with the following steps:

1. An adequate hazard rate function for the individual insurance policies in the index portfolio needs to be chosen. Opting for homogeneity (see third section), we have  $h(t) = \lambda t$  and only the common  $\lambda$  remains unknown.
2. We must estimate the parameters for the jump size distributions, i.e., two for the Gamma model and one each for the chi-squared and the exponential model. Based on their fitted values, we may then directly determine  $\beta$  such that the marginals are kept invariant under the time change.<sup>24</sup>

---

<sup>23</sup>Alternatively, one could have also extracted implied intensity rates separately for each attachment point in a given month. However, this would have resulted in a perfect price fit. Hence, we opt for single cross-section intensities that allows a fair comparison with the estimates obtained by the Lévy-Fraily Model.

<sup>24</sup>See Equation (26).

3. Since the number of constituents in the industry index portfolio and their sums insured are unknown, we do not have an exogenous value for the *pool size*. The same holds true for the recovery rate. Thus, the *adjusted pool size* is endogenously determined in the estimation process according to Equation (30). A drawback is that we are unable to disentangle the pool size from the recovery rate without further information on either of the two.<sup>25</sup>

All three aspects discussed above differ substantially from the analysis of Mai and Scherer (2009a) for structured credit instruments. First, we are unable to derive  $\lambda$  for the marginal distributions from market quotes, since the equivalent of a single-name CDS contract is unavailable in the catastrophe risk space. Second, we explicitly consider three alternative jump size specifications to establish whether the exponential distribution favored by Mai and Scherer (2009a) is suited in the context of ILWs as well. Third, we have to rely on an endogenous estimation of the portfolio size, whereas Mai and Scherer (2009a) calibrate their model to iTraxx Europe tranches, for which the underlying portfolio is known. Hence, adapting the Lévy-Frailty framework to catastrophe risk pricing is not a trivial task.

We fit the models to each end-of-month cross section of ILW quotes in two different ways: through differential evolution developed by Price et al. (2005) and the Nelder-Mead algorithm, also known as direct search. The target criterion is a maximization of the adjusted R-squared ( $\text{adj. } R^2$ ).<sup>26</sup> When an unconstrained optimization is run, both methods yield similar results. Therefore, we decide to report the figures for the much faster direct search approach. For the gamma and exponential specifications, all parameters are estimated simultaneously. In the case of the chi-squared distribution, we proceed slightly differently. Since the degrees of freedom (*df*) parameter  $r$  is an integer, we first define the following fixed range of values:  $r = 1, 2, 3, 4, 5$ . For the subsequent model calibration, we maximize the  $\text{adj. } R^2$  by estimating both the hazard rate parameter  $\lambda$  and the adjusted pool size, while holding  $r$  constant. We separately run this procedure for all of the five values of  $r$ , which results in a set of five  $\text{adj. } R^2$  values. We then select the *df* that produces the largest  $\text{adj. } R^2$ . The optimal value turns out to be  $r = 1$  for both peril types. This result is quite intuitive. For a higher  $r$ , a very large fraction of the pool suffers

---

<sup>25</sup>A more adept treatment of this problem necessitates an empirical analysis of the recovery rate distribution. See Amraoui et al. (2012) for an example in the context of structured credit products.

<sup>26</sup>Note that the optimization requires a numerical implementation of the infinite series in Equation (28). We cut off the series, when the term on the right-hand side of Equation (28) drops below 0.0000001.

impairments if a jump of only average size occurs.<sup>27</sup>

### 4.3 Investigating the Fit

The challenge for a pricing model is to fit all quoted prices with a small set of parameters. Its goodness of fit thus needs to be gauged with respect to error functions. In addition to the adj.  $R^2$ , which was employed for the calibration, We will report the mean absolute error (MAE), the root mean squared error (RMSE) and the mean absolute percentage error (MAPE) associated with the generated prices. These three measures are defined as follows (see, e.g., Armstrong and Collopy, 1992; Xu and Taylor, 1995):

- Root mean square error (RMSE):

$$\text{RMSE} = \sqrt{\frac{1}{N'T'} \sum_{i=1}^{N'} \sum_{t=1}^{T'} (p_{i,t}^{\text{ILW}} - \hat{p}_{i,t}^{\text{ILW}})^2}, \quad (33)$$

- Mean absolute error (MAE):

$$\text{MAE} = \frac{1}{N'T'} \sum_{i=1}^{N'} \sum_{t=1}^{T'} |p_{i,t}^{\text{ILW}} - \hat{p}_{i,t}^{\text{ILW}}|, \quad (34)$$

- Mean absolute percentage error (MAPE):

$$\text{MAPE} = \frac{1}{N'T'} \sum_{i=1}^{N'} \sum_{t=1}^{T'} \left| \frac{p_{i,t}^{\text{ILW}} - \hat{p}_{i,t}^{\text{ILW}}}{p_{i,t}^{\text{ILW}}} \right|. \quad (35)$$

Here,  $p_{i,t}^{\text{ILW}}$  represents the observed market spread for contract  $i$  (characterized by the warranty) at time  $t$ ,  $\hat{p}_{i,t}^{\text{ILW}}$  is the model-generated spread for contract  $i$  at time  $t$ ,  $N' = 7$  denotes the number of contracts for each peril, and  $T' = 97$  equals the number of observations in the time series. Based on these measures, we calculate loss statistics for our entire time-series-cross-section sample of quoted prices. The results for US wind and US earthquake ILWs are depicted in Tables I and II, respectively.<sup>28</sup> In addition,

---

<sup>27</sup>In the case of earthquake contracts, for example, an  $r$  of 2 results in an average trigger probability of 86.47% across all attachment points and months in our data set. This value is derived as follows: we take Equation (28) and compute the trigger probability for each month and all of the seven attachment points. We then average these probabilities per month and subsequently take the mean of all averages over the full time period.

<sup>28</sup>Note that the time series for the calculation of MAPE start on 31 January 2006, since the statistic cannot be calculated

the minimum and the maximum monthly mean absolute error of each model, averaged across the seven contracts, is provided. Apart from the MAPE, the loss measures consistently suggest that all models perform better for US earthquake contracts than for US wind contracts. This finding is owed to the fact that prices and price variations for the former are generally on a lower level (compare Figures 1 and 2).

For wind ILWs (see Table I), the model with chi-squared distributed jump sizes clearly underperforms the other two specifications on all three measures.<sup>29</sup> On its best day (see column “Minimum MAE”), the MAE equals 79 basis points compared to 23 and 29 basis points for the gamma and the exponential model, respectively. Similarly, the MAE, RMSE as well as MAPE are much higher than those of the latter two. Both the gamma and the exponential model achieve an excellent adj.  $R^2$  of 0.93. However, the former additionally exhibits the smallest error statistics across the board. This is not surprising, since the exponential model is nested in the gamma specification. Evidently, the higher precision of the gamma model justifies the additional parameter. The results for the earthquake contracts (see Table II) are very similar. Again, the gamma model dominates the other two approaches and the chi-squared model exhibits the highest errors.

Model	MAE	RMSE	MAPE	Adj. $R^2$
Chi-Squared	0.0107	0.0139	0.0539	0.6247
Exponential	0.0069	0.0045	0.0364	0.9386
Gamma	0.0042	0.0041	0.0299	0.9330
Reduced-Form	0.0579	0.0656	0.4100	0.2951

**Table I: Fit statistics for the time-series-cross-section sample of US wind ILW quotes**

This table shows the goodness of fit for the chi-squared, the exponential and the gamma model, as well as the reduced-form model. All models have been fitted to each end-of-month cross section of US wind ILW quotes between August 2005 and August 2013. Please note that the first three cases use the Nelder-Mead algorithm to maximize the adjusted R-squared (adj.  $R^2$ ), while in the latter case we solve a series of systems of linear equations employing the Levenberg-Marquardt algorithm. As loss statistics, we provide the mean absolute error (MAE), the root mean squared error (RMSE), the mean absolute percentage error (MAPE) and the adjusted R-squared (adj.  $R^2$ ).

---

for longer durations.

<sup>29</sup>A key reason is the boundedness of  $r$  to integer values, which makes this specification less flexible for calibration.

Model	MAE	RMSE	MAPE	Adj. R <sup>2</sup>
Chi-Squared	0.0091	0.0082	0.0710	0.5129
Exponential	0.0048	0.0059	0.0499	0.9049
Gamma	0.0041	0.0049	0.0472	0.8976
Reduced-Form	0.0269	0.0327	0.3548	0.1739

**Table II: Fit statistics for the time-series-cross-section sample of US earthquake ILW quotes**

This table shows the goodness of fit for the chi-squared, the exponential and the gamma model, as well as the reduced-form model. All models have been fitted to each end-of-month cross section of US earthquake ILW quotes between August 2005 and August 2013. Please note that the first three cases use the Nelder-Mead algorithm to maximize the adjusted R-squared (adj. R<sup>2</sup>), while in the latter case we solve a series of systems of linear equations employing the Levenberg-Marquardt algorithm. As loss statistics, we provide the mean absolute error (MAE), the root mean squared error (RMSE), the mean absolute percentage error (MAPE) and the adjusted R-squared (adj. R<sup>2</sup>).

As discussed in the third section, we benchmark the Lévy-Frilty framework with the reduced-form model in the spirit of Jarrow (2010). The corresponding results can be found in the last row of Tables I and II. The pricing errors of this simple but common approach are much higher than for the Lévy-Frilty model with gamma-distributed jump sizes and the adj. R<sup>2</sup> is substantially lower. This holds for both wind as well as earthquake contracts across all four fit statistics and can be traced back to the fact that the model needs to fit each cross-section of monthly ILW spreads with just one parameter.<sup>30</sup> Finally, we put the figures into perspective by comparing them to historical reinsurance quotes.<sup>31</sup> According to data published by Cummins and Danzon (1997), nonproportional property reinsurance prices from different coverage providers exhibited a coefficient of variation of 28 percent between 1980 and 1988.<sup>32</sup> In the same period, the time series of industry prices per unit of ceded exposure was relatively stable (see Froot and O’Connell, 1999). Thus, firm-specific factors are the major driver behind the observed reinsurance price variation (see, e.g., Weiss and Chung, 2004). Against the background of such notable cross-sectional price differences, the MAPE values for the gamma model (2.99 percent and 4.72 percent) reflect a good accuracy.

#### 4.4 Out-of-Sample Analysis

To further assess the suitability of the suggested Lévy-Frilty models for ILW pricing, we run an out-of-sample analysis. As indicated in the first section, the framework is not designed to forecast future ILW prices. Instead, it is meant to be employed on a contemporaneous basis, extracting information

---

<sup>30</sup>Although the model fit is disappointing, unreported results show that the hazard rates implied by the reduced-form model are quite reasonable in economic terms.

<sup>31</sup>Recall from section two that ILWs are essentially a standardized reinsurance contract.

<sup>32</sup>Cummins and Danzon (1997) define prices as premiums earned divided by the present value of accident year losses.



from observed market quotes that can then be harnessed to value natural catastrophe risk transfer instruments for which prices are not observed. Against this background, we adopt the following algorithm for the out-of-sample analysis. Recall that for every month between August 2005 and August 2013, we have two cross sections (one for wind and one for earthquake risk) of seven contracts defined by their attachment points. Hence, in any given month, we fit the model seven times for both cross sections. In each iteration, we reserve a different contract for testing. The model calibrated on the other six contracts is then used to estimate the price of the test contract. In doing so, we pretend that the price of the test contract is unobservable and needs to be estimated from the information included in the prices of the other contracts. Based on this procedure, we are able to calculate an out-of-sample pricing error (estimated minus observed price) for each contract at each point in the time series. These pricing errors are then used to calculate out-of-sample MAE, RMSE, MAPE and  $R^2$ .<sup>33</sup>

The out-of-sample results for all three Lévy-Fraily specifications are summarized in Tables III and IV. As one would expect, the error statistics turn out slightly higher than those measuring the in-sample fit. Importantly though, our general assessment from the previous section is confirmed. The chi-squared model can once more be discarded owing to the fact that it exhibits the worst performance. Moreover, the gamma model outperforms the exponential model, sometimes by a higher margin than in the in-sample analysis. Overall, the gamma model does astonishingly well out-of-sample, indicated by the fact that its error statistics are virtually on par with the in-sample measures for the chi-squared model and far lower than those of the reduced-form model. We may therefore conclude that the gamma model is an adequate choice for the pricing of ILW contracts.

---

<sup>33</sup>Note that the out-of-sample  $R^2$  follows a slightly different logic than the (in-sample) adjusted  $R^2$ . For more information refer, e.g., to Armstrong and Collopy (1992).

Model	MAE	RMSE	MAPE	OOS R <sup>2</sup>
Chi-Squared	0.0486	0.0635	0.2787	0.5606
Exponential	0.0136	0.0505	0.1061	0.8790
Gamma	0.0106	0.0138	0.0677	0.8725

**Table III: Out-of-sample fit statistics for the US wind ILW contracts**

This table shows the goodness of fit for the chi-squared, the exponential and the gamma model, as well as the reduced-form model. All models have been fitted to each end-of-month cross section of US wind ILW quotes between August 2005 and August 2013 seven times, leaving out one of the prices in each iteration. Please note that the first three cases use the Nelder-Mead algorithm to maximize the adjusted R-squared (adj. R<sup>2</sup>), while in the latter case we solve a series of systems of linear equations employing the Levenberg-Marquardt algorithm. The fitted model is then used to predict the price of the contract that was excluded from each estimation. This proceeding allows us to calculate an out-of-sample pricing error (actual minus predicted price) for every contract at every point in time. Based on these pricing errors, we can determine the mean absolute error (MAE), the root mean squared error (RMSE), the mean absolute percentage error (MAPE) and the out-of-sample (OOS) R-squared (adj. R<sup>2</sup>) for each model.

Model	MAE	RMSE	MAPE	OOS R <sup>2</sup>
Chi-Squared	0.0244	0.0340	0.2489	0.4688
Exponential	0.0089	0.0365	0.1198	0.8250
Gamma	0.0075	0.0099	0.1023	0.8305

**Table IV: Out-of-sample fit statistics for the US earthquake ILW contracts**

This table shows the goodness of fit for the chi-squared, the exponential and the gamma model, as well as the reduced-form model. All models have been fitted to each end-of-month cross section of US earthquake ILW quotes between August 2005 and August 2013 seven times, leaving out one of the prices in each iteration. Please note that the first three cases use the Nelder-Mead algorithm to maximize the adjusted R-squared (adj. R<sup>2</sup>), while in the latter case we solve a series of systems of linear equations employing the Levenberg-Marquardt algorithm. The fitted model is then used to predict the price of the contract that was excluded from each estimation. This proceeding allows us to calculate an out-of-sample pricing error (actual minus predicted price) for every contract at every point in time. Based on these pricing errors, we can determine the mean absolute error (MAE), the root mean squared error (RMSE), the mean absolute percentage error (MAPE) and the out-of-sample (OOS) R-squared (adj. R<sup>2</sup>) for each model.

## 4.5 Economic Interpretation

A larger number of parameters generally ensures a higher model flexibility when it comes to adequately fitting observed prices of traded instruments. However, this does not automatically imply that all parameters also have a reasonable economic meaning. Hence, in a last step, we will evaluate the estimated parameter values of the gamma jump-size model with regard to their interpretability. We refrain from a discussion of the other two specifications, since they are clearly dominated in terms of in-sample fit.

According to Tables I and II, the gamma model as the most general specification exhibits the best performance. Taking a closer look at the descriptive statistic for the estimated parameter values in Table V, we find hazard rates and adjusted pool sizes to be relatively stable over time and to move within reasonable bounds. This holds for both peril types. Regarding wind contracts, our estimates for the adjusted pool sizes lie between USD 70.01 bn and USD 331.58 bn. Furthermore, implied hazard rates are found to be in a range from 0.0446 to 0.2144. The adjusted pool sizes for earthquake contracts range

between USD 70.00 bn and USD 521.27 bn, while estimated hazard rates fall between 0.0078 and 0.0999. When investigating the jump frequency that is computed from the shape and the scale parameter, we find it to be of low magnitude and rather constant over time. This is an encouraging result.

The lower bounds of the adjusted pool sizes are just slightly above the highest attachment points of USD 70 bn for wind and USD 50 bn for earthquake contracts. This could be due to relatively low total pool sizes or high recovery values. Since industry index portfolios are arguably larger, we deem the high recoveries to be a more plausible cause for our results. Implied hazard rates, on the other hand, seem quite appropriate, particularly when recalling that they are derived under the risk-neutral pricing measure  $\mathbb{Q}$ . Furthermore, the observation of lower marginal probabilities for earthquake contracts than for wind contracts is consistent with industry experience.

Gamma	Mean	S.D.	Minimum	Median	Maximum
<b>Wind</b>					
Implied Hazard Rate	0.1253	0.0305	0.0446	0.1320	0.2144
Adjusted Pool Size	146.99	48.28	70.01	139.06	331.58
Jump Frequency	3.9317	0.0036	3.1414	3.9364	4.4464
<b>Earthquake</b>					
Implied Hazard Rate	0.0552	0.0203	0.0078	0.0556	0.0999
Adjusted Pool Size	133.60	43.22	70.00	104.34	521.27
Jump Frequency	2.9016	0.0173	2.7199	2.8104	3.2407

**Table V: Descriptive statistics for the estimated parameter values of the gamma model**

This table shows the mean, median, standard deviation (S.D.), as well as the minimum and maximum values for the implied hazard rate ( $\lambda$ ), the adjusted pool size, and the jump frequency of the subordinator ( $\beta$ ) for the model with gamma-distributed jump sizes. The latter has been fitted to each end-of-month cross-section of US wind and US earthquake ILW quotes between August 2005 and August 2013. The adjusted pool size is measured in USD bn.

## 5. Conclusion

We enrich the literature on risk-neutral pricing of ILS by explicitly considering the stochastic dependence of individual claims in the underlying insurance portfolio. Our focus lies on ILWs, since those distinguish themselves through complete standardization and the availability of multiple layers. We highlight why, much like in the case of CDO tranches, the prices for a whole spectrum of different coverage layers must be driven by the stochastic dependence of the policies in the industry index portfolio. Subsequently, we adapt and extend the Lévy-Frailty framework by Mai and Scherer (2009a) to ILWs. More specifically, we employ a compound Poisson subordinator as a stochastic clock that drives loss clustering in the aftermath of catastrophes. We consider three model variants based on chi-squared, exponential and gamma-distributed jump sizes and employ a standard reduced-form model as a benchmark. The Lévy-Frailty framework with gamma-distributed jump sizes is found to exhibit the best fit to a comprehensive set of historical ILW prices, both in sample and out of sample.

Several directions for future research are conceivable. First, our analysis could be enriched with data on the claim intensities of the individual policies. Such information should be available in insurance companies and could be used to further improve the model fit by explicitly estimating the common hazard rate. Second, it would also be desirable to disentangle the recovery value from the pool size through an exogenous determination of those variables. The former could be based on an empirical distribution of the percentage loss on the policies, given a claim event has occurred, and the latter should be obtainable from an insurance industry index provider such as PCS. Third, although we focused on ILWs, the approach could be extended to promote the market-consistent valuation of catastrophe risk in general.

## References

- Amraoui, S., Cousot, L., Hitier, S., and Laurent, J.-P. (2012). Pricing CDOs with state-dependent stochastic recovery rates. *Quantitative Finance*, 12(8):1219–1240.
- Ané, T. and Geman, H. (2000). Order flow, transaction clock, and normality of asset returns. *The Journal of Finance*, 55(5):2259–2284.
- Aon Benfield (2019). Reinsurance market outlook January 2019. Technical report.
- Armstrong, S. J. and Collopy, F. (1992). Error measures for generalizing about forecasting methods: Empirical comparisons. *International Journal of Forecasting*, 08(1):69–80.
- Bakshi, G. and Madan, D. (2002). Average rate claims with emphasis on catastrophe loss options. *Journal of Financial and Quantitative Analysis*, 37(01):93–115.
- Balbás, A., Longarela, I. R., and Lucia, J. J. (1999). How financial theory applies to catastrophe-linked derivatives. An empirical test of several pricing models. *Journal of Risk and Insurance*, 66(4):551–581.
- Bantwal, V. J. and Kunreuther, H. C. (2000). A cat bond premium puzzle? *Journal of Psychology and Financial Markets*, 1(1):76–91.
- Beer, S. and Braun, A. (2019). Market-consistent valuation of natural catastrophe risk.
- Borovkova, S. and Schmeck, M. D. S. (2017). Electricity price modeling with stochastic time change. *Energy Economics*, 63:51 – 65.
- Böttcher, B. (2010). Feller processes: The next generation in modeling. Brownian motion, Lévy processes and beyond. *PLoS One* 5, e15102, 5(12):1–8.
- Braun, A. (2011). Pricing catastrophe swaps: A contingent claims approach. *Insurance: Mathematics and Economics*, 49(3):520–536.
- 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., Affolter, S., and Schmeiser, H. (2016). Life settlement funds: Current valuation practices and areas for improvement. *Risk Management and Insurance Review*, 19(2):173–195.

- Braun, A., Ben Ammar, S., and Eling, M. (2019a). Asset pricing and extreme events: Common factors in ILS fund returns. *Journal of Banking and Finance*, 102(5):59–78.
- Braun, A., Gatzert, N., and Schmeiser, H. (2012). Performance and risks of open-end life settlement funds. *Journal of Risk and Insurance*, 79(1):193–230.
- Braun, A., Luca, D., and Schmeiser, H. (2019b). Consumption-based asset pricing in insurance markets: Yet another puzzle? *Journal of Risk and Insurance*, 86(3):629–661.
- Braun, A., Müller, K., and Schmeiser, H. (2013). What drives insurers' demand for cat bond investments? evidence from a pan-european survey. *The Geneva Papers on Risk and Insurance - Issues and Practice*, 38(3):580–611.
- Braun, A., Schmeiser, H., and Schreiber, F. (2015). Solvency ii's market risk standard formula: How credible is the proclaimed ruin probability? *Journal of Insurance Issues*, 38(1):1–30.
- Braun, A., Schmeiser, H., and Schreiber, F. (2017). Portfolio optimization under solvency ii: Implicit constraints imposed by the market risk standard formula. *Journal of Risk and Insurance*, 84(1):177–207.
- Braun, A., Schmeiser, H., and Schreiber, F. (2018). Return on risk-adjusted capital under solvency ii: Implications for the asset management of insurance companies. *The Geneva Papers on Risk and Insurance - Issues and Practice*, 43(3):456–472.
- Braun, A. and Weber, J. (2017). Evolution or revolution? how solvency ii will change the balance between reinsurance and ils. *Journal of Insurance Regulation*, 36(4):1 – 26.
- Carr, P. and Wu, L. (2004). Time-changed Lévy processes and option pricing. *Journal of Financial Economics*, 71(1):113–141.
- Chang, C. W., Chang, J. S. K., and Lu, W. (2008). Pricing catastrophe options in discrete operational time. *Insurance: Mathematics and Economics*, 43(3):422–430.
- Chang, C. W., Chang, J. S. K., and Lu, W. (2010). Pricing catastrophe options with stochastic claim arrival intensity in claim time. *Journal of Banking and Finance*, 34(1):24–32.

- Chang, C. W., Chang, J. S. K., and Yu, M.-T. (1996). Pricing catastrophe insurance futures call spreads: A randomized operational time approach. *Journal of Risk and Insurance*, 63(4):599–617.
- Clark, P. K. (1973). A subordinated stochastic process model with finite variance for speculative prices. *Econometrica: Journal of the Econometric Society*, pages 135–155.
- Cummins, J. and Danzon, P. M. (1997). Price, financial quality, and capital flows in insurance markets. *Journal of Financial Intermediation*, 6(1):3 – 38.
- Cummins, J. D. et al. (2006). Should the government provide insurance for catastrophes? *Federal Reserve Bank of St. Louis Review*, 88(4):337–379.
- Cummins, J. D. and Weiss, M. A. (2009a). Convergence of insurance and financial markets: Hybrid and securitized risk-transfer solutions. *The Journal of Risk and Insurance*, 76(3):493–545.
- Cummins, J. D. and Weiss, M. A. (2009b). Convergence of insurance and financial markets: Hybrid and securitized risk-transfer solutions. *Journal of Risk and Insurance*, 76(3):493–545.
- Embrechts, P. and Meister, S. (1997). Pricing insurance derivatives: The case of cat-futures. In *Proceedings of the 1995 Bowles Symposium on Securitization of Insurance Risk, Georgia State University Atlanta*, ed. by S. Cox, pages 15–26.
- Froot, K. and O’Connell, P. (1999). *The Pricing of US Catastrophe Reinsurance*, page 195–232. University of Chicago Press, Chicago.
- Gatzert, N. and Schmeiser, H. (2012). Industry loss warranties: Contract features, pricing, and central demand factors. *The Journal of Risk Finance*, 13(1):13–31.
- Gatzert, N., Schmeiser, H., and Toplek, D. (2011). An analysis of pricing and basis risk for industry loss warranties. *Zeitschrift für die gesamte Versicherungswissenschaft*, 100(4):517–537.
- Gatzert, N., Sebastian, P., and Nikolai, V. (2017). Convergence of capital and insurance markets: Consistent pricing of index-linked catastrophe loss instruments. *Journal of Risk and Insurance*, forthcoming.
- Geman, H. (2002). Pure jump Lévy processes for asset price modelling. *Journal of Banking and Finance*, 26(7):1297–1316.

- Geman, H. and Yor, M. (1997). Stochastic time changes in catastrophe option pricing. *Insurance: Mathematics and Economics*, 21(3):185–193.
- Giesecke, K. (2003). A simple exponential model for dependent defaults. *The Journal of Fixed Income*, 13(3):74–83.
- Guy Carpenter (2016). Catastrophe bond update: Fourth quarter and full year 2015. Technical report.
- Jaffee, D. M. and Russell, T. (1997). Catastrophe insurance, capital markets, and uninsurable risks. *Journal of Risk and Insurance*, pages 205–230.
- Jarrow, R. A. (2010). A simple robust model for cat bond valuation. *Finance Research Letters*, 7(2):72–79.
- Jarrow, R. A. and Turnbull, S. M. (1995). Pricing derivatives on financial securities subject to credit risk. *The Journal of Finance*, 50(1):53–85.
- Joshi, M. S. and Stacey, A. M. (2006). Intensity gamma: A new approach to pricing portfolio credit derivatives. Technical report.
- Kalemanova, A., Schmid, B., and Werner, R. (2007). The normal inverse Gaussian distribution for synthetic CDO pricing. *The Journal of Derivatives*, 14(3):80–94.
- Kielholz, W. and Durrer, A. (1997). Insurance derivatives and securitization: New hedging perspectives for the US cat insurance market. *Geneva Papers on Risk and Insurance. Issues and Practice*, pages 3–16.
- Lindskog, F. and McNeil, A. J. (2003). Common Poisson shock models: Applications to insurance and credit risk modelling. *Astin Bulletin*, 33(02):209–238.
- Mai, J.-F. (2010). *Extendibility of Marshall-Olkin distributions via Lévy subordinators and an application to portfolio credit risk*. Dissertation.
- Mai, J.-F. and Scherer, M. (2009a). A tractable multivariate default model based on a stochastic time-change. *International Journal of Theoretical and Applied Finance*, 12(02):227–249.
- Mai, J.-F. and Scherer, M. (2009b). Lévy-frailty copulas. *Journal of Multivariate Analysis*, 100(7):1567–1585.



- Marshall, A. W. and Olkin, I. (1967). A multivariate exponential distribution. *Journal of the American Statistical Association*, 62(317):30–44.
- McDonnell, E. (2002). Industry loss warranties. In Lane, M., editor, *Alternative Risk Strategies*, chapter 4, pages 81–98. Risk Books, a division of the Risk Water Group.
- NOAA (2011). Hurricanes and tropical storms for annual 2010. Technical report, National Oceanic and Atmospheric Administration.
- O’Kane, D. and Turnbull, S. (2003). Valuation of credit default swaps. *Lehman Brothers Fixed Income Quantitative Credit Research*.
- PartnersRe (2015). The drivers of catastrophe bond pricing. Technical report.
- Price, K. V., Storn, R. M., and Lampinen, J. A. (2005). The differential evolution algorithm. *Differential Evolution: A Practical Approach to Global Optimization*, pages 37–134.
- Tankov, P. (2003). *Financial modelling with jump processes*, volume 2. CRC press.
- Weiss, M. and Chung, J.-H. (2004). U.s. reinsurance prices, financial quality, and global capacity. *Journal of Risk and Insurance*, 71(3):437–467.
- Xu, X. and Taylor, S. J. (1995). Conditional volatility and the information efficiency of the PHLX currency options market. *Journal of Banking and Finance*, 19(5):803–821.
- Zeng, L. (2000). On the basis risk of industry loss warranties. *The Journal of Risk Finance*, 1(4):27–32.