



HAL
open science

Combination of traditional and parametric insurance: calibration method based on the optimization of a criterion adapted to heavy tail losses

Olivier Lopez, Daniel Nkamani

► **To cite this version:**

Olivier Lopez, Daniel Nkamani. Combination of traditional and parametric insurance: calibration method based on the optimization of a criterion adapted to heavy tail losses. 2025. <hal-04959706v3>

HAL Id: hal-04959706

<https://hal.science/hal-04959706v3>

Preprint submitted on 23 Jul 2025 (v3), last revised 16 Feb 2026 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC0 1.0 - Universal - International License

Combination of traditional and parametric insurance: calibration method based on the optimization of a criterion adapted to heavy tail losses.

Olivier LOPEZ¹, Daniel NKAMENI^{1,2}

July 23, 2025

Abstract

In this paper, we consider the question of providing insurance protection against heavy tail losses, where the expectation of the loss may not even be finite. The product we study is based on a combination of traditional insurance up to some limit, and a parametric (or index-based) cover for larger losses. This second part of the cover is computed from covariates available just after the claim, allowing to reduce the claim management costs via an instant compensation. To optimize the design of this second part of the product, we use a criterion which is adapted to extreme losses (that is distribution of the losses that are of Pareto type). We support the calibration procedure by theoretical results that show its convergence rate, and empirical results from a simulation study and a real data analysis on tornados in the US. We conclude our study by empirically demonstrating that the proposed hybrid contract outperforms a traditional capped indemnity contract.

Key words: Parametric insurance; heavy-tailed distributions; natural disasters; capped indemnity contracts.

Short title: Combination of Traditional and Parametric insurance for heavy-tailed losses.

¹ CREST Laboratory, CNRS, Groupe des Écoles Nationales d'Économie et Statistique, Ecole Polytechnique, Institut Polytechnique de Paris, 5 avenue Henry Le Chatelier 91120 PALAISEAU, France

² Detralytics, 1-7 Cours Valmy 92923 Paris La Défense Cedex, France. E-mails: olivier.lopez@ensae.fr, daniel.nkameni@ensae.fr

1 Introduction

The modern landscape of insurance is characterized by the emergence of new risks whose potential of destruction question the feasibility of developing a sustainable financial protection. Several reports¹ alerted about the potential uninsurability of natural disasters in the context of climate change. In the context of cyber, the difficulty to cover risks that are increasing in terms of frequency, severity and variety is materialized by significant protection gaps². To overcome these challenges, parametric insurance (or "index-based insurance") is often mentioned as one of the technical tools that may help offering a coverage in these apparently desperate situations: a report from the French Direction of Treasury on cyber insurance³ mention these "innovative products" as a potential solution to respond to the need for financial protection against this expanding threat.

The concept of parametric insurance offers indeed a convenient solution to the difficulty to cover risks that are difficult to apprehend, either because they are too volatile, or because lack of data makes it hard to develop models required to build an efficient and profitable protection. The trick is essentially not to cover the risk itself, but to build a product for which compensation is determined by the values taken by a parameter (or index) available just after the claim. Multiple examples of application can be found for example in Abdi et al. [2022], Eze et al. [2020], Lin and Kwon [2020], Carter et al. [2017] or Alderman and Haque [2007]. Of course, the compensation mechanism has to be determined so that it matches with the true loss experienced by the policyholder, that is one needs to control the "basis risk" (see Doherty and Richter [2002] or Elabed et al. [2013]). But if this difference between the loss and the compensation is contained, the product benefits from many advantages:

- the parameter is usually considered because many data are available on it and/or because its distribution is easier to control from a risk management perspective than the one of the true loss. Better knowledge of the parameter means more control on the compensation mechanism.
- the availability of the parameter reduces expertise and claim management costs, it also ensures a faster compensation for the policyholder.

¹see the report for the European Insurance and Occupational Pensions Authority, https://www.eiopa.europa.eu/leveraging-insurance-shore-europes-climate-resilience-2024-09-03_en or the report on insurability of climate risk mandated by the French ministry of Economy, see https://www.ecologie.gouv.fr/sites/default/files/documents/Rapport_final_Mission-assurance_climat.pdf.

²See for example the report from the Global Federation of Insurance Associations in March 2023, <https://gfaiinsurance.org/topics/487>.

³see <https://www.tresor.economie.gouv.fr/Articles/2022/09/07/remise-du-rapport-sur-le-developpement-de-l-assurance-du-risque-cyber>.

- the rules are clear, and this should lead to fewer legal dispute between policyholders and insurers.

Another advantage of this type of mechanism is the possibility to develop Insurance Linked Securities that are based on these parametric products, allowing to transfer the risk to capital markets, see Braun et al. [2023], Securities [2009], Michel-Kerjan and Morlaye [2008]. Although we will not explore this dimension in the present paper, focusing on the ability of such a product to cover a damage, this dimension should not be forgotten as an important leverage to mobilize more capital (via the markets) in the absorption of catastrophic risks, making these products more appealing in a broader perspective.

In the present paper, we consider the design of parametric insurance as a complement of traditional insurance, when the distribution of the loss is heavy-tailed. In the field on Insurance Linked Securities, Härdle and Cabrera [2010] or Mastroeni et al. [2022] are examples where a Pareto-type variable models the underlying risk. In this case, full coverage of the loss by a traditional insurance contract may be difficult due to a too important premium, or even impossible if the expectation of the loss is infinite. The parametric insurance product is used on top of a limited insurance contract to improve the protection at an affordable price. Many market study show the popularity of parametric solutions to cover the upper segment of the risk⁴.

To design optimal parametric insurance, several approaches have been studied in the literature. A first category includes optimization of criteria related to utility theory, as in Louaas and Picard [2021], or Tan and Zhang [2024], with various elaborate techniques have been proposed, from regression trees to neural networks in Chen et al. [2024]. In presence of heavy tail risk, Conradt et al. [2015] relied on quantile regression, see also Gu et al. [2023] in the case of earthquakes. Our approach differs in the sense that we consider a different metric, focusing on the difference between the loss and the compensation offered by the product, but rather on the ratio. This framework seems adapted to the study of heavy tail distributions, first of all because it avoids dealing with quantities that may have infinite expectation. Also because, in this context, the risk is so important that it is implicit that it is impossible to cover it perfectly. On the other hand, the percentage of the loss that is left to the policyholder is a concrete indicator of the performance of the solution, on which it is easy to communicate.

Summary of contributions

The contributions of this paper are driven by the need to determine the optimal protection mechanism of the policyholder, when this mechanism is made of a combination of traditional insurance, supplemented by some parametric product to cover the upper segment. Our analysis is performed in four steps:

⁴see for example the report from PricewaterhouseCoopers, <https://www.pwc.ch/en/insights/fs/closing-the-gap-with-parametrics-insurance.html>

1. We first define in section 2 a function that we want to optimize, in conceiving the product. The idea is to find a compromise between the proportion of the loss that is covered, and an affordable price.
2. Once this metric has been defined, we define an optimal cover based on the available information after the claim (see section 2.3 to 2.5).
3. We provide a statistical methodology to approach this optimal policy based on available data (section 3).
4. We investigate the behavior of the procedure on simulations (section 4) and illustrate the feasibility of the method, on tornado loss data from the US (section 5).
5. We compare the proposed contract to the capped indemnity contract which is generally used by insurers (section 6)

The technical proofs are postponed to section 8.

2 Index based reinsurance

This section is devoted to defining the framework of the insurance cover that we study in this paper, namely a combination of "traditional" and parametric insurance. The first step is to introduce, in section 2.1, a proper metric, adapted to measure what is expected from this product, in the context of heavy tail losses. A part of this metric is made to a way to materialize the aversion to high prices, which is discussed in section 2.2. We then describe a specific type of product based on a combination of traditional insurance with parametric insurance taking over after some threshold in section 2.3. A first way to design the index-based part of the product is discussed in section 2.4.

2.1 A metric to optimize

In the following, we consider a policyholder who will be affected by the loss Y (which may be zero if no claims happen during the period under study). In this section, we propose a way to measure how an insurance contract responds to the needs of the policyholder, in the particular case where Y is heavy tail, that is

$$S_Y(t) = \mathbb{P}(Y \geq t) = \frac{l(t)}{t^{1/\gamma}}, \quad (2.1)$$

where $\gamma > 0$ is called the tail index and l is a slow-varying function, that is

$$\forall h > 0, \lim_{t \rightarrow \infty} \frac{l(ht)}{l(t)} = 1. \quad (2.2)$$

This class of distributions plays an important role in extreme value analysis, see Beirlant et al. [2004], and is particularly adapted to consider risk with huge potential losses.

Consider an insurance contract that offers the compensation X , usually smaller than Y . We want to maximize the proportion of the risk that is covered, that is the ratio X/Y . In case of large risks, this indicator can be more convenient than looking at the difference between X and Y : if the tail of Y is heavier than the tail of X , this may exhibit large gaps that blurry the perception of the performance. The ratio is less affected by this phenomenon. Hence, in the following, we consider that the aim is to maximize

$$\mathfrak{L}_X = E \left[L \left(\frac{X}{Y} - f(\pi_X) \right) \right], \quad (2.3)$$

where L is a non-decreasing function, f is a non-decreasing function that materializes the aversion to high prices, and with the convention $0/0 = 1$.

For the function L , an obvious choice is the identity $L(x) = x$, but other choices could be more suitable as discussed in section 3.1. The role of the function f is discussed in section 2.2.

2.2 Price aversion

We now discuss the choice of the function f . In (2.3), the benefits of an increase of the ratio X/Y can be compensated by an increase of the price, hence of $f(\pi_X)$. Since the ratio X/Y is supposed to be bounded⁵, it is natural to consider a non-decreasing function $f : \mathbb{R} \rightarrow [0, 1]$ that is bounded, like, for example,

$$f_l(\pi) = \frac{\kappa}{1 + \exp(-\beta\pi)},$$

and a rational fraction

$$f_r(\pi) = \frac{\kappa\pi^\beta}{1 + \pi^\beta}.$$

From a statistical perspective, the function f plays the role of a penalty function in the selection of the best cover X among a given family. In this vision, the quantity κ in f_r and f_l can be understood as a hyperparameter that controls the force of the penalization.

Let us discuss some implications of the choice of a specific form of f . Consider the special case where, $X = X(s) = \min(Y, s)$. We have

$$\pi(s) = E[\min(Y, s)] = \int_0^s S_Y(t) dt.$$

Moreover, the function (2.3) (that we here denote $\mathfrak{L}(s)$ to shorten the notation) becomes

$$\mathfrak{L}(s) = (1 - S_Y(s))L(1 - f(\pi(s))) - \int_s^\infty L\left(\frac{s}{t} - f(\pi(s))\right) dS_Y(t). \quad (2.4)$$

⁵The spirit is to have $X/Y \leq 1$. Overcompensation can exist, but should be particularly rare, otherwise the product would not be interesting anymore for the insurer. Therefore X/Y is expected to be bounded by a constant slightly larger than 1 except of an event a small probability.

The derivative of (2.4) with respect to s is

$$\mathfrak{L}'(s) = -\pi'(s)f'(\pi(s))(1-S_Y(s))L'(1-f(\pi(s))) - \int_s^\infty L'\left(\frac{s}{t} - f(\pi(s))\right) \left[\frac{1}{t} - \pi'(s)f'(\pi(s))\right] dS_Y(t).$$

Let us consider a few particular situations:

1. $E[Y] = \pi^+ < \infty$. When s tends to infinity, $\pi(s) \rightarrow \pi^+$ and $\pi'(s) = S(s) \rightarrow 0$. If $f'(\pi^+) > 0$, the first term is

$$-\frac{l(s)f'(\pi^+)}{s^{1/\gamma}}L'(1-f(\pi^+)) + o\left(\frac{1}{s^{1/\gamma}}\right),$$

while the second is positive but is $O(s^{-1/\gamma-1})$. Hence, for s large enough, $\mathfrak{L}'(s) < 0$, which means that it corresponds to a situation where the policyholder is not necessarily seeking for a perfect protection against the risk (which would be $s = \infty$): after some point, the burden of a high premium is too important, and \mathfrak{L} decreases.

2. $E[Y] = \infty$. Then, if $f(\pi(s))$ is bounded, $f'(\pi(s)) \rightarrow 0$ as s tends to infinity. In this case, $\mathfrak{L}'(s)$ is equivalent, for large s , to $-l(s)f'(\pi(s))s^{-1/\gamma}L'(1-f(\pi(s)))$ only if gets to zero faster than $s^{-1/\gamma-1}$ (the positive term in the decomposition of \mathfrak{L}'), which happens if $l(s)sf'(\pi(s)) \rightarrow \infty$. If this condition does not hold, this corresponds to the situation where the policyholders are ready to pay any amount of premium against a perfect protection. Note that condition $l(s)sf'(\pi(s)) \rightarrow \infty$ holds for function f_r if $\beta < \gamma/(\gamma-1)$ in the case $\gamma > 1$, since $\pi(s) = O(s^{1-1/\gamma})$ and $f_r(\pi) = O(\pi^{-(1+\beta)})$ when π tends to infinity. On the other hand, it is not the case for f_l since $f_l(\pi) = O(\exp(-\beta\pi))$.

Clearly, the practical choice of f should be calibrated from data describing the behavior of the policyholders, but this discussion shows that some shapes can be eliminated based on some basic knowledge on their behavior regarding the possibility of full protection (ready to pay any price or not).

We now decline the general framework (2.3) in the particular case of the parametric insurance protection that we aim to study.

2.3 Risk and index based compensation

In the following, we consider a loss variable Y with survival function $S_Y(t) = \mathbb{P}(Y \geq t)$. The policyholder wants protection against this loss for an affordable price. On the other hand, the insurance coverage is partial, and the pay-off of the insurance contract is

$$X_\phi(s) = Y\mathbf{1}_{Y \leq s} + s\phi(\mathbf{W})\mathbf{1}_{Y > s}, \quad (2.5)$$

with $\mathbf{W} \in \mathbb{R}^d$ is a set of variables that are observed soon after the claim, and will allow to compute a compensation $s\phi(\mathbf{W})$. The spirit is that $\phi(\mathbf{W})$ should, in most cases, be larger than

1, but it could be authorized to have $\phi(\mathbf{W})$ less than 1 to "smooth" the behavior around s and to avoid over-compensation (see also section 2.5). Moreover, the resulting under-compensation could be tolerated by the policyholder due to the fast payment associated with index-based products. The threshold s is essentially a way for the insurer to protect himself against the too large values taken by Y . But instead of simply giving a limit of compensation s , once this threshold is reached, a new mechanism of compensation is used, based on \mathbf{W} .

This framework is designed to mimic the behavior of a parametric (or "index-based") insurance products that completes the traditional cover. Here we made a simplification: the parametric part of the product, $s\phi(\mathbf{W})$, is triggered by a condition on the value of Y . The behavior of parametric insurance products can be considered in two steps:

- a triggering mechanism, namely a condition for the compensation to activate;
- a compensation, if activated.

In practice, both of these components should depend on observable variables \mathbf{W} that are available soon after the occurrence of the claim, this is the idea of fastness of compensation related to parametric insurance. This is not the case here, since we assume that the triggering mechanism is based on a decision based on the value of Y ($\{Y > s\}$), which is a quantity known after a significant amount of time (since it requires expertise). We introduce this simplification since the focus of this paper is more on the second part of the process, that is how to design the compensation once the guarantee has been activated. In practice, the events $\{Y \leq s\}$ and $\{Y > s\}$ should be replaced by events of the type $\{\tilde{\phi}(\mathbf{W}) \leq s\}$ and $\{\tilde{\phi}(\mathbf{W}) > s\}$ for some function $\tilde{\phi}$ such that $\mathbb{P}(\{Y > s\} \cap \{\tilde{\phi}(\mathbf{W}) \leq s\}) + \mathbb{P}(\{Y \leq s\} \cap \{\tilde{\phi}(\mathbf{W}) > s\})$ is close to zero. We will discuss how this extension slightly modifies our results in section 2.5.

We can see the product X_ϕ as a kind of reinsurance mechanism, where s would represent the priority of a stop-loss reinsurance treaty. The quantity $s\phi(\mathbf{W})$ is expected to be close to Y , and, essentially, $s\phi(\mathbf{W}) \leq Y$ (at least with high probability) otherwise the compensation would exceed the true loss (a financial product could be designed with such pay-off, but this is not the spirit of insurance products which only aim at repairing and not at generating profit for the policyholder). The pure premium corresponding to this contract is denoted $\pi_\phi^*(s) = E[X_\phi(s)]$. It is assumed to be defined, that is $\pi^*(s) < \infty$, meaning that $E[X] < \infty$ even in the situation where $E[Y] = \infty$. Moreover, qualitatively, we should keep in mind that this quantity is less than the pure premium (if defined) of an insurance contract integrally covering Y , otherwise this insurance mechanism lacks interest. The final premium $\pi_\phi(s)$ paid by the policyholder includes some loading factor τ , that is

$$\pi_\phi(s) = (1 + \tau)\pi_\phi^*(s).$$

Let us note that here and in the rest of the paper, the frequency of occurrence of the claims is

taken into account implicitly: the variable Y may take the value 0 with a non zero probability, materializing the fact that a proportion of policyholders will not experience any claim.

2.4 A first method of selection of an optimal payoff function ϕ

The question is now to find the appropriate payoff function ϕ to maximize the criterion (2.3). This is not solely a matter of approximating the true loss Y : if it was the case, and if $E[Y^2]$ were finite, a possible way to proceed would be to consider $\phi(\mathbf{w}) = s^{-1}E[Y|Y \geq s, \mathbf{W} = \mathbf{w}]$, which would minimize the mean square error between $X_\phi(s)$ and Y . Here, quadratic loss is not the proper metric for heavy tail variables (that may even have infinite expectation and/or variance). Moreover, this would not respect the aversion to high prices materialized by the function ϕ .

To select the proper function ϕ , let us consider that we have at our disposal a class of potential functions $\mathcal{F} = \{\phi_\theta; \theta \in \Theta\}$, with $\Theta \subset \mathbb{R}^k$. The dimension k may be very large, for example if one relies on a over-parametrized machine learning model.

In Louaas and Picard [2021] and Chen et al. [2024], the authors propose to maximize an empirical version of the objective function to select the appropriate value of θ (hence the corresponding function ϕ_θ). Transposed in our context, and assuming that we have i.i.d. historical data $(Y_i, \mathbf{W}_i)_{1 \leq i \leq n}$, the idea would be to maximize over Θ the quantity, leading to

$$\hat{\theta}_n(s) = \arg \max_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n L \left(\frac{X_{\theta,i}(s)}{Y_i} - f(\hat{\pi}_\theta(s)) \right), \quad (2.6)$$

where, to shorten the notation, we use $X_{\theta,i}(s)$ (resp. $\pi_\theta(s)$) instead of $X_{\phi_\theta,i}(s) = Y_i \mathbf{1}_{Y_i \leq s} + s \phi_\theta(\mathbf{W}_i) \mathbf{1}_{Y_i > s}$ (resp. $\pi_{X_{\phi_\theta}(s)}$), $\hat{\pi}_\theta(s) = (1 + \tau)n^{-1} \sum_{i=1}^n X_{\theta,i}(s)$.

Let us note that, if the dimension of Θ is small compared to n , which will be the case in a classical parametric model, under standard assumptions for M-estimators (see Van Der Vaart et al. [1996]), we get $\hat{\theta}_n(s) - \tilde{\theta}(s) = O_P(n^{-1/2})$, where

$$\tilde{\theta}(s) = \arg \max_{\theta \in \Theta} E \left[L \left(\frac{X_\theta(s)}{Y} - f(\pi_\theta(s)) \right) \right] \quad (2.7)$$

corresponds to the parameter of the best payoff function from the class \mathcal{F} . This rate of convergence is informative on the volume of historical data required to approximate the ideal payoff. One of the practical difficulties stands in the fact that the size of this database can be limited, especially for emerging risks: this may be hard to collect a large amount of data where Y_i and \mathbf{W}_i are simultaneously present. On the other hand, it is much easier to get additional data on \mathbf{W} , since the measure of these variables used to compute the index can be obtained without any expensive and/or time-consuming expertise.

Therefore, based on a closer analysis of (2.3), we will recommend a two-step procedures that can, in some cases, increase the convergence rate, as described in section 3 below.

2.5 Trigger mechanism based on the covariates

In parametric insurance, the idea is not to rely on an exact knowledge of Y to engage the compensation. In this section, we study the case where the trigger mechanism is of the type $\{s\phi(\theta, \mathbf{W}) > \tilde{s}\}$, instead of $\{Y > s\}$, and where the compensation, for the traditional part of the product, is limited by s .

Let

$$X^*(s, \tilde{s}) = Y \mathbf{1}_{s\phi(\tilde{\theta}(s), \mathbf{W}) \leq \tilde{s}} + s\phi_{\tilde{\theta}(s)}(\mathbf{W}) \mathbf{1}_{s\phi_{\tilde{\theta}(s)}(\mathbf{W}) > \tilde{s}},$$

where $\tilde{\theta}(s)$ has been defined in (2.7). The idea is to have something that (under appropriate conditions) is close to the product $X_{\tilde{\theta}(s)}(s)$, without needing to observe Y to trigger the parametric part of the product.

The product defined by $X_{\tilde{\theta}(s)}(s)$, has somehow an optimal trigger condition: it ensures that all claims with amount Y less than s are completely compensated, and that there are no claims with amount larger than s that require claim management (except for determining the exact amount of the claim). To approximate this ideal trigger mechanism, the idea is that $\{s\phi(\theta, \mathbf{W}) > \tilde{s}\}$, should be designed to approximate the event $\{Y > s\}$. Two types of mismatch can of course occur in two ways:

- $Y > s$ and the parametric part of the protection is not activated. In this case, the policyholder will be fully compensated, and this compensation will exceed the level s . This is a problem for the insurer which has to deal with a large claim that exceeds the limit, but the index part of the contract is not used. The idea is to take \tilde{s} small enough (less than s) so that this situation does not occur too often.
- $Y \leq s$ and $s\phi_{\tilde{\theta}(s)}(\mathbf{W}) > \tilde{s}$. In this case, the policyholder can be disadvantaged if Y is between s and \tilde{s} , and if $Y \geq s\phi_{\tilde{\theta}(s)}(\mathbf{W})$. Although it is not taken into account in this paper, in this case, $s\phi_{\tilde{\theta}(s)}(\mathbf{W})$ stays close to Y and $s - \tilde{s}$ is small, the fast payment in the case of parametric insurance could compensate this small loss for the policyholder. On the other hand, if $Y \leq \tilde{s}$, there is overcompensation, which should still be avoided.

Proposition 2.1 shows that there are few differences between the products $X_{\tilde{\theta}(s)}(s)$ and $X^*(s, \tilde{s})$ in terms of compensation under some appropriate condition. This allows, in the optimization of the payoff, to use the simple strategy (2.6) (or the alternative strategy defined in the next section), rather than optimizing the value of θ with respect to the role it plays in the trigger mechanism.

Proposition 2.1 *Assume that L' and f' are bounded. Let*

$$\begin{aligned} E_-(s, \tilde{s}) &= \left\{ Y \leq s \text{ and } s\phi_{\tilde{\theta}(s)} \geq \tilde{s} \right\}, \\ \Delta(s, \tilde{s}) &= E \left[\left(s\phi_{\tilde{\theta}(s)}(\mathbf{W}) - Y \right)_+ \mathbf{1}_{E_-(s, \tilde{s})} \right]. \end{aligned}$$

and let $\pi^*(s, \tilde{s}) = (1 + \tau)E[X^*(s, \tilde{s})]$, $\tilde{\pi}(s) = (1 + \tau)E[X_{\tilde{\theta}(s)}(s)]$. Then,

$$E \left[L \left(\frac{X^*(s, \tilde{s})}{Y} - f(\pi^*(s, \tilde{s})) \right) \right] \geq E \left[L \left(\frac{X_{\tilde{\theta}(s)}(s)}{Y} - f(\tilde{\pi}(s)) \right) \right] - C(\mathbb{P}(E_-(s, \tilde{s})) + \Delta(s, \tilde{s})),$$

for some constant $C > 0$.

The lower bound obtained in this result shows that if we optimize the value of θ in X_θ , we can use the same value for a product whose trigger mechanism does not depend on Y (but can be computed from \mathbf{W}): in this case, the objective functions (\mathfrak{L}) of both products is almost the same, up to some perturbations that depend on the quality of the trigger mechanism. The proof is provided in section 9. The interpretation of the quantity $\Delta(s, \tilde{s})$ is related to the potential risk of overcompensation of the index-based part of the product, that should stay controlled as explained above.

3 A new metric to optimize an index-based product

In the section, we consider a collection of covers

$$X_\theta(s) = Y\mathbf{1}_{Y \leq s} + s\phi_\theta(\mathbf{W})\mathbf{1}_{Y > s},$$

where we want to select, from data, the best function ϕ_θ from a collection \mathcal{F} , using the notations of section 2.4.

Our first task is to obtain an approximation of

$$\mathfrak{L}_\theta(s) = E \left[L \left(\frac{X_\theta(s)}{Y} - f(\pi_\theta(s)) \right) \right],$$

which is done in section 3.1. From this approximation, we deduce in section 3.2 an alternative way to determine a approximate the ideal cover, that is the quantity $\tilde{\theta}(s)$ defined by (2.7). The benefits of relying on this alternative to the approximation $\hat{\theta}(s)$ defined in (2.6) are shown in section 3.3 where we study the convergence rates of the different procedures.

3.1 Approximation of the criterion

Our first step is to investigate the behavior of $\mathfrak{L}_\theta(s)$ when s is large. The need for this approximation comes from the fact that we want to design an index-based cover which completes the coverage of the policyholder only for large claims. The design of this parametric part of the product depends on analyzing the conditional distribution of Y given \mathbf{W} where we recall that \mathbf{W} are covariates that are available just after the claim. Modeling the whole distribution of $Y|\mathbf{W}$ can be relatively complex, and the idea is that our final result should be essentially influenced by the tail distribution. Getting rid of the remaining part of the distribution is a simplification,

avoiding to make too strong assumptions. This is the spirit of (2.1), where the distribution of Y is specified only in terms of rate of decay of S_Y , without specifying completely the function l .

Before stating Theorem 3.1, which is the main result of this section, we need some generic assumption on the distribution of $Y|\mathbf{W}$ and the functions L and f involved in the computation of $\mathfrak{L}_\theta(s)$.

Assumption 1 *Let $S(t|\mathbf{w}) = \mathbb{P}(Y \geq t|\mathbf{W} = \mathbf{w})$. Assume that, for all w ,*

$$S(t|\mathbf{w}) = \frac{l(t|\mathbf{w})}{t^{1/\gamma(\mathbf{w})}},$$

where $\gamma(\mathbf{w}) > 0$, and $l(\cdot|\mathbf{w})$ is a slow-varying function. Moreover, assume that, for some constants C and $c > 0$, there exists a slow varying function \mathfrak{l} such that

$$cl(t) \leq l(t|\mathbf{w}) \leq C\mathfrak{l}(t).$$

Moreover we assume that

$$\sup_{\theta, \mathbf{w}} \left| \frac{l(st|\mathbf{w})}{l(s|\mathbf{w})} \right| \rightarrow 1,$$

when s tends to infinity.

Assumption 1 above states that the conditional distribution of Y remains in the same family (2.1), but with index parameter potentially depending on \mathbf{W} . The introduction of \mathfrak{l} is essentially here to allow us to apply the dominated convergence theorem, controlling the fact that there is no structural change in the conditional distribution of Y .

Next, we need two assumptions on the function L and f .

Assumption 2 *Let $\pi^+ = \lim_{s \rightarrow \infty} \pi_\theta(s)$. There exists a function $\varphi_0(t) > 0$ such that*

$$\sup_t \left| \frac{\frac{L(t-f(\pi))}{L(1-f(\pi))} - \varphi_0(t)}{\varphi_0(t)} \right| \rightarrow_{\pi \rightarrow \pi^+} 0.$$

Assumption 3 *There exists a function $\varphi_1(t) < 0$ such that*

$$\sup_t \left| \frac{\frac{L'(t-f(\pi))}{L(1-f(\pi))} - \varphi_1(t)}{\varphi_1(t)} \right| \rightarrow_{\pi \rightarrow \pi^+} 0..$$

Regarding the definition of the function L ,

1. If $L(x) = L_1(x) = x$, then

$$\begin{aligned} \varphi_0(t) &= \frac{t - f(\pi^+)}{1 - f(\pi^+)}, \\ \varphi_1(t) &= \frac{1}{1 - f(\pi^+)}, \end{aligned}$$

under the condition that $f(\pi^+) < 1$.

Another definition of L could be :

2. $L_2(x) = -\exp(-\mu x)$. In this case,

$$\begin{aligned}\varphi_0(t) &= \exp(-\mu(t-1)), \\ \varphi_1(t) &= -\mu \exp(-\mu(t-1)).\end{aligned}$$

L_2 introduces non-linearity and concavity into the resulting optimization problem. This has the advantage of stabilizing the optimization and numerical results, and also facilitates the identification of the global optimum. Moreover, L_2 is more robust to outliers and large values—an especially useful property when dealing with extreme losses.

These advantages of L_2 over L_1 can be explained by the fact that L_2 leads to an optimization that implicitly models the trade-off between mean and variability, unlike L_1 , which only targets the mean. Indeed, if X_θ is a random variable, applying a Taylor expansion around $\mathbf{E}[X]$ gives:

$$\mathbf{E}[-\exp(-\mu X_\theta)] \approx -\exp\left(-\mu \mathbf{E}[X_\theta] + \log\left\{1 + \frac{\mu^2}{2} \mathbf{V}[X_\theta]\right\}\right)$$

Thus, L_2 approximates a penalty on the variance \mathbf{V} .

Finally, through a Taylor expansion, L_2 can also be interpreted as a generalization of L_1 that incorporates higher-order powers of the variable of interest.

Theorem 3.1 *Under Assumption 1 to 3,*

$$\mathfrak{L}_\theta(s) = L(1 - f(\pi_\theta(s)))E[1 - S(s|\mathbf{W})\Phi_0(\phi(\theta, \mathbf{W}), \gamma(\mathbf{W}))](1 + \mathfrak{R}_\theta(s)),$$

where

$$\begin{aligned}\Phi_1(x, \gamma) &= \int_0^x v^{1/\gamma} \varphi_1(v) dv, \\ \Phi_0(x, \gamma) &= 1 - \varphi_0(x) + \Phi_1(x, \gamma),\end{aligned}$$

with $\sup_\theta |\mathfrak{R}_\theta(s)| \rightarrow 0$ when s tends to infinity.

The proof is postponed to section 9.1. We will use the notation $\mathfrak{L}_\theta^*(s)$ to refer to the main term in the expansion of Theorem 3.1, that is

$$\mathfrak{L}_\theta^*(s) = \mathfrak{L}_\theta(s)/(1 + \mathfrak{R}_\theta(s)).$$

A first remark is that

$$\mathfrak{L}_\theta^*(s) = E[\Psi(\pi_\theta(s), \phi(\theta, \mathbf{W}); S(s|\mathbf{W}), \gamma(\mathbf{W}))]. \quad (3.1)$$

One can observe, in $\mathfrak{L}_\theta^*(s)$, a separation between a term $L(1 - f(\pi_\theta(s)))$ that materializes the effect of the price, and the term between brackets. The two terms drive \mathfrak{L} in two opposite directions.

Another interesting feature of this decomposition is that we see that the whole joint distribution of Y and \mathbf{W} is not involved in Ψ : to compute the approximation of $\mathfrak{L}(s)$, one only needs to compute the conditional tail index $\mathbf{w} \rightarrow \gamma(\mathbf{w})$ and $\mathbf{w} \rightarrow S(s|\mathbf{w})$, and to know the distribution of \mathbf{W} . Essentially, one needs to estimate the probability of triggering the parametric compensation $s\phi(\theta, \mathbf{w})$ (which is easier to do than if we needed to estimate the whole conditional distribution of Y above s), and since the information \mathbf{W} is usually more available than precise claim data on Y , this distribution can be estimated with a good accuracy (see section 3.2 below).

3.2 Alternative optimization for the parametric part

Let us consider that we have two databases. The first contains joint i.i.d. replications $(Y_i, \mathbf{W}_i)_{1 \leq i \leq n}$ of (Y, \mathbf{W}) . The second contains i.i.d. replications $(\mathbf{W}_j)_{1 \leq j \leq m}$ of \mathbf{W} (the first n observations are common to the two databases, although there is no obligation).

Theorem 3.1 encourages to rely on a two step procedure to optimize the index based cover:

- Step 1: from $(Y_i, \mathbf{W}_i)_{1 \leq i \leq n}$, estimate $w \rightarrow S(s|\mathbf{w})$ and $\mathbf{w} \rightarrow \gamma(\mathbf{w})$. Let $\hat{S}(s|\mathbf{w})$ and $\hat{\gamma}(\mathbf{w})$ denote the corresponding estimators.
- Step 2: Let us recall that $\hat{\pi}_\theta(s) = (1 + \tau)n^{-1} \sum_{i=1}^n X_{\theta,i}(s)$. Let

$$\hat{\mathfrak{L}}_\theta^*(s) = \frac{1}{m} \sum_{j=1}^m \Psi \left(\hat{\pi}_\theta(s), \phi(\theta, \mathbf{W}_j); \hat{S}(s|\mathbf{W}_j), \hat{\gamma}(\mathbf{W}_j) \right),$$

where Ψ is defined in (3.1). Then, define

$$\hat{\theta}^*(s) = \arg \max_{\theta \in \Theta} \hat{\mathfrak{L}}_\theta^*(s). \quad (3.2)$$

The heuristic is the following: data on both Y and \mathbf{W} are rare (and potentially expensive, see André et al. [2013]). They are used to provide estimators of $\hat{S}(s|\cdot)$ and $\hat{\gamma}(\cdot)$. The convergence rate of these estimators will be limited by the sample size n . On the other hand, we have plenty of information on the covariates \mathbf{W} , which means that m is large. The errors on \hat{S} and $\hat{\gamma}$ are averaged in (3.2), which increases the performance. This improvement will be limited, first by the bias introduced by the fact that $\hat{\mathfrak{L}}_\theta^*(s)$ is an empirical approximation of $\mathfrak{L}_\theta^*(s)$ instead of $\mathfrak{L}_\theta(s)$. This bias is supposed to be controlled by the fact that s should be selected high enough. A second limitation will become apparent with the conditions of section 3.3, but the rate of convergence of $\hat{\theta}^*(s)$ will be, at worse, the one of $\hat{S}(s|\mathbf{w})$ and $\hat{\gamma}(\mathbf{w})$.

Regarding possible estimators of $\hat{S}(s|\cdot)$ and $\hat{\gamma}(\cdot)$, examples will be provided in section 5. A logistic regression can for example be used for $\hat{S}(s|\cdot)$. However, tail regression techniques may be used to simultaneously estimate this probability and the parameter $\hat{\gamma}$, see for example Farkas et al. [2021] and Farkas et al. [2024] who rely regression trees. The next section is then devoted to providing theoretical ground to legitimate the use of (3.2).

3.3 Convergence rate for the statistical optimization

We here give theoretical elements to support the convergence of $\hat{\theta}^*(s)$ towards $\tilde{\theta}^*(s)$, where

$$\tilde{\theta}^*(s) = \arg \max_{\theta \in \Theta} \mathfrak{L}_\theta^*(s).$$

Let us note that $\tilde{\theta}^*(s)$ is different from $\tilde{\theta}(s)$ due to the approximation that is done in $\mathfrak{L}_\theta^*(s)$, but Theorem 3.1 shows that this bias can be contained if s is large enough. This convergence requires a set of standard assumptions which are the adaptation of standard in (semi-)parametric M-estimation. Essentially, they require:

- regularity of the functions $\theta \rightarrow \mathfrak{L}_\theta^*(s)$, which is essentially a matter of regularity of the function L and f , see Assumption 4 below;
- the estimation technique used for \hat{S} and $\hat{\gamma}$ should be consistent and not too "unstable". This is the meaning of the more technical Assumptions 5 and 6, see also the discussion below.

Let us also stress that we are focusing on a situation where k , the dimension of Θ , is small compared to n and m , to rely on asymptotic theory (for n and m tending to infinity with n potentially much lower than m). Hyper-parametrized techniques like deep neural networks are clearly beyond the scope of these theoretical results.

We now list the assumptions required for Theorem 3.2.

Assumption 4 *We assume the following conditions:*

1. *the functions $|L|$, $|L'|$, $|f'|$ and φ_0 and φ_1 are bounded;*
2. *there exists a constant M_0 such that*

$$\sup_{\theta \in \Theta} \left| \frac{\partial \pi_\theta(s)}{\partial \theta} \right| \leq M_0;$$

3. *there exists a function Λ with $E[\Lambda(\mathbf{W})] < \infty$ such that*

$$\left| \frac{\partial \phi(\theta, \mathbf{w})}{\partial \theta} \right| \leq \Lambda(\mathbf{w}).$$

4. *we have $\gamma_- = \inf_{\mathbf{w} \in \mathcal{W}} \gamma(\mathbf{w}) \geq \frac{1}{\alpha}$, for some $\alpha > 0$, with*

$$\int_0^\infty v^\alpha \varphi_1(v) dv < \infty.$$

One can link this assumption with Theorem 3.2.5 in Van Der Vaart et al. [1996]. The idea is essentially to ensure that, if we knew exactly S and γ , the empirical M-estimation procedure would be $m^{1/2}$ -consistent.

Assumption 5 Let \mathcal{W} be such that $\mathbb{P}(\mathbf{W} \in \mathcal{W}) = 1$. We have

$$\sup_{\mathbf{w} \in \mathcal{W}} \left| \frac{\hat{S}(s|\mathbf{w}) - S(s|\mathbf{w})}{1 - S(s|\mathbf{w})} \right| + \sup_{\mathbf{w} \in \mathcal{W}} |\hat{\gamma}(\mathbf{w}) - \gamma(\mathbf{w})| = O_P(\varepsilon_n),$$

with ε_n tending to zero.

Assumption 5 is related to the consistency of the preliminary method of estimation of the conditional tail distribution of Y . Convergence rates of this type can be found in Farkas et al. [2024] for the case of estimation via regression trees.

Assumption 6 is more technical, and is related to the concept of Donsker class, see Van Der Vaart et al. [1996], Chapter 2.1 for a definition.

Assumption 6 1. The functions $\hat{S}(s|\mathbf{w})$ (resp. $\mathbf{w} \rightarrow \hat{\gamma}(\mathbf{w})$ and $\mathbf{w} \rightarrow S(s|\mathbf{w})$ (resp. $\mathbf{w} \rightarrow \hat{\gamma}(\mathbf{w})$) belong to a Donsker class of functions \mathcal{H}_1 (resp. \mathcal{H}_2). Moreover \mathcal{H}_1 is bounded.

2. The classes of functions $\{\mathbf{w} \rightarrow \phi(\theta, \mathbf{w}) : \theta \in \Theta\}$ and $\{\mathbf{w} \rightarrow \partial_\theta \phi(\theta, \mathbf{w}) / \partial \theta : \theta \in \Theta\}$ are Donsker. Moreover the second class is bounded.

Basically, it ensures that the estimators $\hat{S}(s|\mathbf{w})$ and $\hat{\gamma}(\mathbf{w})$ do not evolve in too complex classes of functions (point 1), that would be too large for a uniform central limit property to hold. Typically, it will hold for differentiable functions (with bounded derivatives up to some order depending on the dimension of \mathbf{w} , see for example Theorem 2.7.1 in Van Der Vaart et al. [1996]), or piecewise constant functions. For point 2, it will essentially hold if $\phi(\theta, \mathbf{w})$ and its gradient with respect to θ are Lipschitz with respect to θ (see Theorem 2.7.11 in Van Der Vaart et al. [1996]), and if the dimension of Θ can be considered as negligible with respect to m .

We now state the main result of this section.

Theorem 3.2 Let

$$B_n = \int \nabla_\theta \Psi \left(\pi_{\tilde{\theta}^*}(s), \phi(\tilde{\theta}^*(s), \mathbf{w}); \hat{S}(s|\mathbf{w}), \hat{\gamma}(\mathbf{w})) \right) d\mathbb{P}(\mathbf{w}),$$

and

$$\Sigma(s) = \nabla_\theta^2 \mathcal{L}_{\tilde{\theta}^*}^*(s),$$

which is assumed to be invertible. Under the Assumptions of Theorem 3.1 and Assumptions 4 to 6,

$$\hat{\tilde{\theta}}^*(s) - \tilde{\theta}^*(s) = -\Sigma(s)^{-1} \left\{ \frac{1}{m} \sum_{j=1}^m \nabla_\theta \Psi \left(\pi_{\tilde{\theta}^*}(s), \phi(\tilde{\theta}^*(s), \mathbf{W}_j); S(s|\mathbf{W}_j), \gamma(\mathbf{W}_j)) \right) \right\} + O_P(B_n) + O_P(m^{-1/2}).$$

The proof is provided in section 9.2.

In this decomposition, the empirical sum is tending to 0 at rate $m^{-1/2}$. The term B_n is tending to zero, since, from the definition of $\tilde{\theta}^*(s)$,

$$\int \nabla_{\theta} \Psi \left(\pi_{\tilde{\theta}^*}(s), \phi(\tilde{\theta}^*(s), \mathbf{w}); S(s|\mathbf{w}), \gamma(\mathbf{w})) \right) d\mathbb{P}(\mathbf{w}) = 0,$$

and \hat{S} and $\hat{\gamma}$ converge towards S and γ . The rate convergence will be at worse ε_n , but if $\nabla_{\theta} \Psi \left(\pi_{\tilde{\theta}^*}(s), \phi(\tilde{\theta}^*(s), \mathbf{w}); \hat{S}(s|\mathbf{w}), \hat{\gamma}(\mathbf{w})) \right)$ is an unbiased estimator of its limit, $B_n = 0$ thanks to the integration with respect to $\mathbb{P}(\mathbf{w})$ (and if the bias is small, it can lead to a significant gain). This property will be confirmed by the empirical example studied in the next section.

4 Empirical analysis on simulated data

In this section, we will illustrate the theoretical results obtained so far on simulated data. The work in this paper applies to heavy-tailed losses, which are the focus of all the analyses in this section. Section 4.1 presents the simulation setting, the preliminary manipulations of the data, and the analysis algorithm used. Section 4.2 presents the results obtained from these analyses.

4.1 Setting

These analyses have two main objectives. The first is to perform an empirical comparison of the objective function \mathfrak{L}_{θ} and its approximation \mathfrak{L}_{θ}^* on empirical data for a given threshold s . The second is to empirically verify that a two-step estimation of the objective function, using the approximation \mathfrak{L}_{θ}^* and an influx of index data, could help improve the quality of the estimation.

For these analyses, we choose a framework with a one-dimensional index \mathbf{W} and a heavy-tailed loss Y . \mathbf{W} is assumed to follow a uniform distribution on $[0, 1]$. Y is assumed to depend on the index \mathbf{W} and to follow a Pareto distribution whose survival function is given by:

$$S(Y = t|\mathbf{W} = \mathbf{w}) = \mathbb{P}(Y > t|\mathbf{W} = \mathbf{w}) = \frac{1}{t^{\frac{1}{\gamma(\mathbf{w})}}}.$$

The link between Y and \mathbf{W} is established through the shape parameter $\gamma(\mathbf{W})$, which is defined as:

$$\gamma(\mathbf{w}) = e^{-a-b\mathbf{w}}, \tag{4.1}$$

where a and b are parameters chosen such that $\gamma(\mathbf{W}) \in [0.2, 0.7]$ for all $W \in [0, 1]$.

The threshold s is chosen as the 85th percentile of the losses Y . In other words, the proposed index reinsurance is applied to the top 15% of losses. Moreover, the loading factor τ is 10%. The family of payoff functions $\phi_{\theta}(\mathbf{W})$ is defined as:

$$\phi_{\theta}(\mathbf{w}) = \frac{1}{s} \max\{\min[\mathbf{E}(Y|Y > s, \mathbf{W} = \mathbf{w}), e^{\theta\mathbf{w}}], s\}.$$

In this expression, the minimum between $\mathbf{E}(Y|Y > s, \mathbf{W} = \mathbf{w})$ and $e^{\theta \mathbf{w}}$ is used to ensure there is no overcompensation, at least on average. In practice, this adjustment could be modified if it results in an unacceptable proportion of under-compensation. Such modifications could involve adding a loading parameter to $\mathbf{E}(Y|Y > s, \mathbf{W} = \mathbf{w})$ or adapting the variable part $e^{\theta \mathbf{w}}$ and the domain of θ to implicitly account for and prevent overcompensation. The maximum between the previous expression and s is used to guarantee that there is no under-compensation above s . This ensures that a policyholder who suffers a loss greater than s receives at least s as compensation.

The optimal value of θ is therefore obtained under the constraint that, for losses above s , the policyholder or the cedent (in the case of index reinsurance) receives compensation of at least s and at most $\mathbf{E}(Y|Y > s, \mathbf{W} = \mathbf{w})$ (which is finite since $\gamma(\mathbf{w})$ stays below 1).

We consider $L(x) = L_2(x) = -\exp(-\mu x)$, where μ is set to 1.5. For the price aversion function, we consider $f = f_r$, where f_r is as defined in section 2.2, with $\kappa = 1.415$ and $\beta = 1.65$. Note that, in practice, these values would be calibrated using information from the portfolio and the policyholders.

4.2 Results

Figure 1 presents an analysis of the objective function, \mathfrak{L} , and its approximation \mathfrak{L}^* , on the full sample of size $m = 5000$. Plots of \mathfrak{L} and \mathfrak{L}^* are shown in panel (a). These curves will serve as references in our subsequent analyses. We observe that the curves of \mathfrak{L} and \mathfrak{L}^* are nearly identical, which empirically confirms the validity of the using \mathfrak{L}^* instead of \mathfrak{L} . The slight difference between the two curves can be attributed mainly to the error introduced by the approximation of the integral present in the expression of \mathfrak{L}^* (see Theorem 3.1).

Panel (b) shows the payoff function ϕ_θ for the optimal value of θ as a function of the actual losses of the policyholders. The ideal payout function, which corresponds to the exact loss Y , is also plotted. We observe that the limits s and $\mathbf{E}(Y|Y > s, \mathbf{W} = \mathbf{w})$ are quite effective, as they seem to successfully prevent undercompensation and overcompensation, respectively. The optimal value of θ , which maximizes the empirical objective functions, produces a payout close to the ideal payout. This is evidenced by the clustering of the payouts above s around the line $\phi_\theta = Y$.

The deviations from the ideal payout are not necessarily problematic, as they will be accounted for in the price of the insurance contract, ensuring that policyholders pay exactly for the amount of coverage they receive. Of course, this is guaranteed only if the index model is sufficiently accurate.

The next part of our analysis consists of studying and illustrating the advantages of a two-step optimization of θ using \mathfrak{L}^* . For this, according to the setting of section 3.2, we consider a

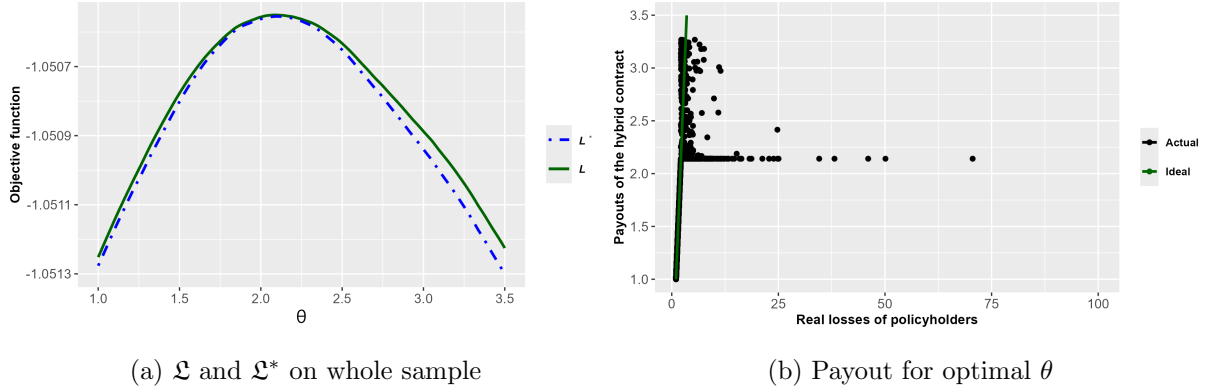


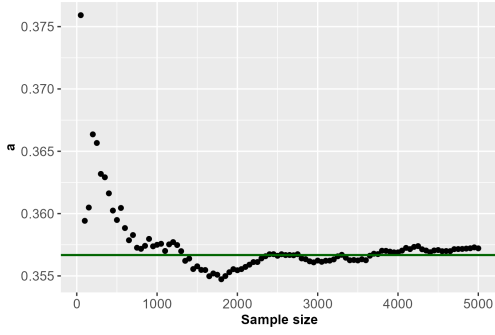
Figure 1: Panel (a) shows the objective functions for the whole sample of size m as a function of the parameter θ for both \mathcal{L} and \mathcal{L}^* . Panel (b) shows the payout function $\phi_{\hat{\theta}}$ as a function of the real losses of policyholders. $\hat{\theta}$ is the optimal value of θ obtained from the maximization of $\hat{\mathcal{L}}_{\theta}$.

situation where we have a sample $(Y_i, \mathbf{W}_i)_{1 \leq i \leq n}$ of size n (where we will make this size n vary) and a larger sample $(\mathbf{W}_j)_{1 \leq j \leq m}$ of size $m > n$. Algorithm 1 outlines the steps followed in this analysis.

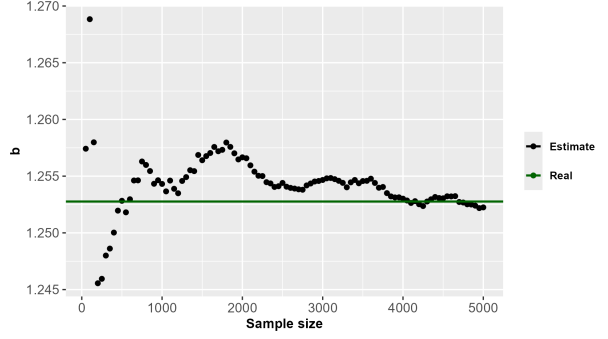
Algorithm 1 Methodology of analysis

- 1: $m \leftarrow$ total sample size
 - 2: $n \leftarrow$ a starting sample size
 - 3: $k \leftarrow$ an increment
 - 4: $\hat{\mathcal{L}}_{ref}(\theta) \leftarrow \frac{1}{m} \sum_{i=1}^m \left[L \left(\frac{Y_i \mathbf{1}_{Y_i \leq s} + s \phi_{\theta}(W_i) \mathbf{1}_{Y_i > s}}{Y_i} - f(\hat{\pi}_{\theta}(s)) \right) \right]$
 - 5: **repeat**
 - 6: $\hat{\mathcal{L}}_n(\theta) \leftarrow \frac{1}{n} \sum_{i=1}^n \left[L \left(\frac{Y_i \mathbf{1}_{Y_i \leq s} + s \phi_{\theta}(\mathbf{W}_i) \mathbf{1}_{Y_i > s}}{Y_i} - f(\hat{\pi}_{\theta}(s)) \right) \right]$
 - 7: Estimate \hat{a}_n and \hat{b}_n using $(Y_i, W_i)_{1 \leq i \leq n}$ to obtain $\hat{\gamma}(W)$ and $\hat{S}(s|W)$
 - 8: $\hat{\mathcal{L}}_n^*(\theta) \leftarrow L (1 - f(\hat{\pi}_{\theta}(s))) \frac{1}{m} \sum_{i=1}^m \left[1 - \hat{S}(s|W_i) \Phi_0(\phi_{\theta}(W_i), \hat{\gamma}(W_i)) \right]$
 - 9: $Error_n \leftarrow \|\hat{\mathcal{L}}_n, \hat{\mathcal{L}}_{ref}\|_{\infty} = \sup_{\theta} |\hat{\mathcal{L}}_n(\theta) - \hat{\mathcal{L}}_{ref}(\theta)|$
 - 10: $Error_n^* \leftarrow \|\hat{\mathcal{L}}_n^*, \hat{\mathcal{L}}_{ref}\|_{\infty} = \sup_{\theta} |\hat{\mathcal{L}}_n^*(\theta) - \hat{\mathcal{L}}_{ref}(\theta)|$
 - 11: store $Error_n$, $Error_n^*$ and the parameters \hat{a}_n and \hat{b}_n
 - 12: $n \leftarrow n + k$
 - 13: **until** $n = m$
-

In step 7 of this algorithm, we use the smaller sample $(Y_i, \mathbf{W}_i)_{1 \leq i \leq n}$ to estimate the parameters a and b of equation 4.1. These estimated parameters are sufficient to define the functions $\mathbf{w} \mapsto \gamma(\mathbf{w})$ and $\mathbf{w} \mapsto S(s|\mathbf{w})$ from (4.1) and the definition of $S(t|\mathbf{w})$ from section 4.1. In step 8, these two functions are incorporated into the expression of \mathcal{L}^* , which does not depend on



(a) Estimates of the parameter a

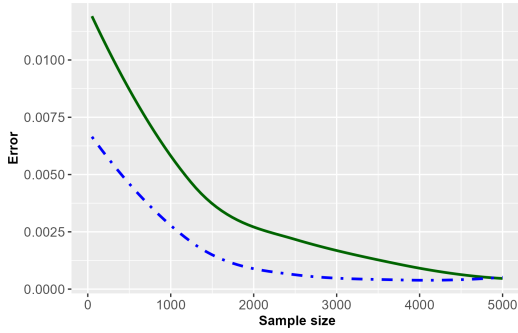


(b) Estimates of the parameter b

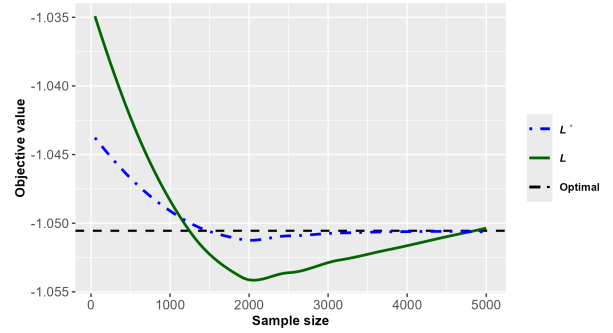
Figure 2: Evolution of the estimates of the parameters a (panel (a)) and b (panel (b)) of equation 4.1 with sample size.

Y or its entire distribution. An estimate $\hat{\mathcal{L}}^*$ of \mathcal{L}^* is then calculated using the larger sample $(\mathbf{W}_i)_{1 \leq i \leq m}$.

Figure 2 illustrates the evolution of the estimates of the parameters a and b as a function of the sample size n . As expected, as the size of the sample used for estimation increases, the estimates converge to the true values of these parameters. These true values, which serve as references, are estimated using the full sample of size m .



(a) Errors in estimating \mathcal{L} by \mathcal{L}_n and \mathcal{L}_n^* .



(b) Optimal values of \mathcal{L}_n , \mathcal{L}_n^* and \mathcal{L} .

Figure 3: Panel (a) shows the errors incurred in the estimation of \mathcal{L} by \mathcal{L}_n (in one step) and \mathcal{L}_n^* (in two steps) as a function of n (see steps 7 and 8 of algorithm 1). Panel (b) shows the objective values obtained through \mathcal{L}_m , \mathcal{L}_n , and \mathcal{L}_n^* . The evolution of the latter two is presented as a function of n .

The results obtained by applying algorithm 1 are presented in figure 3. In panel (a), we observe a decrease in the estimation errors of \mathcal{L} by \mathcal{L}_n (one-step estimation) and \mathcal{L}_n^* (two-step estimation) as the sample size of $(Y_i, \mathbf{W}_i)_{1 \leq i \leq n}$ increases. We also observe that the two-step estimation of the objective function using \mathcal{L}_n^* results in a lower estimation error compared to the

one-step estimation using \mathfrak{L}_n . This result supports the conclusion that in situations with limited information on economic losses, the proposed methodology can improve the estimation of the optimal index insurance model parameters θ by leveraging a larger sample of index values.

In practice, beyond addressing the scarcity of economic loss data for emerging risks, this approach could also help in accounting for the future evolution of index values in current estimations, provided reliable projections of index values are available. Such an approach could significantly enhance the risk management strategies of index insurers and reinsurers.

The efficiency of the two-step approach using \mathfrak{L}^* is further corroborated by panel (b) of figure 3. Specifically, the objective values calculated in two steps using \mathfrak{L}_n^* are closer to, and converge faster toward, the reference objective value compared to the objective values calculated in one step using \mathfrak{L}_n .

In addition to reducing the number of random variables and the amount of information required to estimate the objective function \mathfrak{L}^* , there may be an additional practical advantage to using \mathfrak{L}^* . In a highly favorable scenario where the insurance or reinsurance company possesses a very large dataset on Y and \mathbf{W} , estimating and optimizing \mathfrak{L} directly on this large dataset could be computationally expensive. In such a case, an alternative solution could involve estimating $\hat{\gamma}(\mathbf{W})$ and $\hat{S}(s|\mathbf{W})$ on a smaller subset of the dataset, incorporating these estimates into \mathfrak{L}^* , and then performing the optimization using only \mathbf{W} .

5 Empirical analysis on real data

In this section, we present the results of our analysis using real data. We consider tornado loss data from the United States covering the period 2016 to 2023. Section 5.1 describes this dataset and the setting used for our analysis, while section 5.2 presents the results obtained.

5.1 Data and setting

The tornado losses dataset⁶ used in this study is identical to the one analyzed by Daouia et al. [2023]. In this work, the authors confirmed the heavy-tailed nature of these losses over the period 2010 to 2016 by fitting a Generalized Pareto Distribution.

The variables of interest in this dataset include tornado monetary losses (*loss*), the starting and ending longitude and latitude of each tornado (*slon*, *elon*, *slat*, and *elat*), and the length and width of the area affected by the tornado (*len* and *wid*). To address issues related to the scale of the losses, we work with losses normalized per unit area, defined as:

$$Y = \frac{loss}{len \times wid}.$$

⁶Available at <https://www.spc.noaa.gov/wcm/#/data>

Similarly to Daouia et al. [2023], we chose the tornado’s average geographical position as indices, defined as:

$$W^{(1)} = \frac{slat + elat}{2}, \text{ and } W^{(2)} = \frac{slon + elon}{2},$$

and consider $\mathbf{W} = (W^{(1)}, W^{(2)})$. Including additional indices and covariates could make sense in the practical design of an index insurance or reinsurance product for tornado losses. However, due to the lack of data and to ensure the simplicity of our analysis, we limit ourselves to this two dimensional vector of covariates.

Our final sample $(Y_i, \mathbf{W}_i)_{1 \leq i \leq m}$ contains $m = 4659$ observations and covers the entire USA. In particular, our dataset includes the costliest tornado, which occurred on October 20, 2019, causing EF3 damage in the Dallas suburbs of Texas and resulting in \$1.5 billion in losses. Let us note that, in this example, we do not have two samples, one with (Y, \mathbf{W}) and one with \mathbf{W} only, but we will artificially remove some values of Y_i to measure the performance of the approximation \mathcal{L}^* .

Figure 4 presents the tornado losses per unit area over the study period. We distinguish between losses below the 85th percentile, which corresponds to our choice of s , and losses above this threshold. Despite a concentration of losses in the eastern regions, extreme tornadoes (losses above the 85th percentile) occur across nearly all states in the US. This observation leads to the broader conclusion that no geographical location is entirely immune to extreme losses.

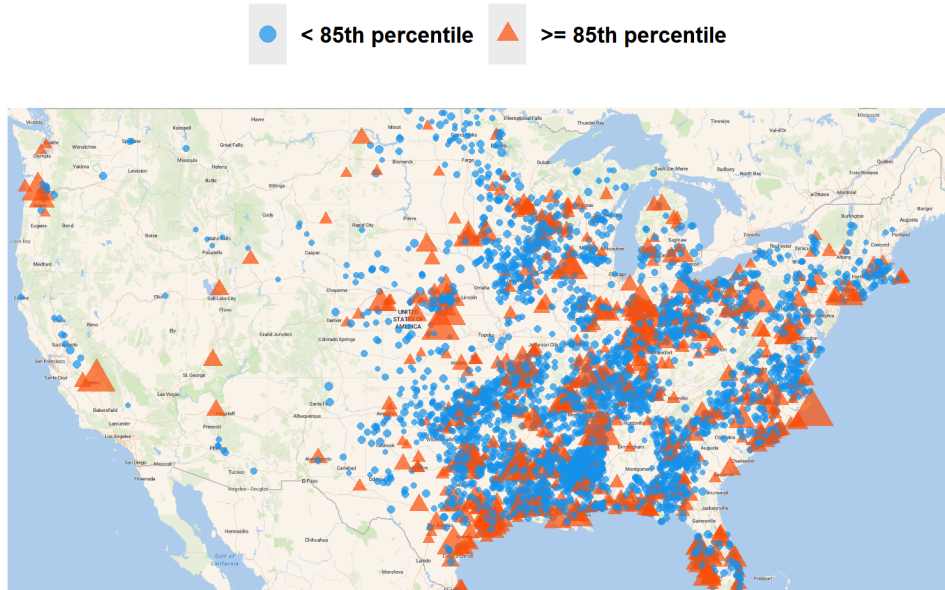


Figure 4: Tornado losses per unit area at the starting locations of the tornadoes. A distinction is made between losses per unit area below and above the 85th percentile, which is the value chosen for s .

We assume a Generalized Pareto Distribution (GPD) for the tornado losses per unit area, with a dependence between these losses and the index values, such that:

$$S(Y = t | \mathbf{W} = \mathbf{w}) = \frac{1}{\left(1 + \frac{y\gamma(\mathbf{w})}{\sigma}\right)^{\frac{1}{\gamma(\mathbf{w})}}}.$$

The shape parameter γ , which establishes the statistical link between the losses Y and the indices $W^{(1)}$ and $W^{(2)}$, is defined as:

$$\gamma(\mathbf{w}) = e^{-a-bw^{(1)}-cw^{(2)}}, \quad (5.1)$$

where a , b , and c are parameters to be estimated from the data. As mentioned earlier, s is chosen as the 85th percentile of the tornado losses per unit area. In practice, other choices of s could be made depending on the setting and context without hindering the applicability of the proposed methodology.

The family of payoff functions is defined as:

$$\phi_{\theta_1, \theta_2}(\mathbf{w}) = \frac{1}{s} \max\{\min[q_{0.5}(Y|Y > s, \mathbf{W} = \mathbf{w}), \theta^{(1)}w^{(1)} + \theta^{(2)}w^{(2)}], s\}$$

where $q_{0.5}(Y|Y > s, \mathbf{W} = \mathbf{w})$ is the conditional median of losses per unit area above s . This statistic replaces the conditional expectation used for simulated data in section 4.1 because, for certain tornado locations, the shape parameter is greater than 1, rendering the conditional expectation undefined.

The choice of the median could be adapted in practice depending on the structure of the index model and the distribution of losses, as long as the selected statistic prevents overcompensation and limits under-compensation to an acceptable level. The justification for this choice of structure for the payout function is similar to that provided in section 4.1 for simulated data.

The values of μ , the parameters of the function f , and the loading factor of the premium are identical to those used for the simulated data.

5.2 Results

Panel (a) of Figure 5 shows a plot of \mathcal{L}^* with respect to θ_1 and θ_2 for $s = q_{0.85}(Y)$. We observe a strict concavity of \mathcal{L}^* with respect to θ_1 and θ_2 , suggesting the existence of an optimal combination (θ_1, θ_2) that maximizes the objective function. Panel (b) of this figure shows plots of \mathcal{L}^* for values of s in $q_{0.84}(Y)$, $q_{0.85}(Y)$, $q_{0.86}(Y)$. We observe that as the value of the quantile increases, meaning as we focus on more extreme losses, the value of the objective function appears to increase.

This behavior is due to the implicit assumption that all losses below s are fully and efficiently covered by indemnity-based insurance at the same cost as index insurance. In practice, this is

not always the case, as indemnity-based insurance is not perfectly accurate. Furthermore, if we account for delays in compensation and high costs associated with indemnity-based insurance for extreme losses, it is reasonable to expect that, in practice, indemnity insurance would not always be preferred by policyholders for high or extreme losses.

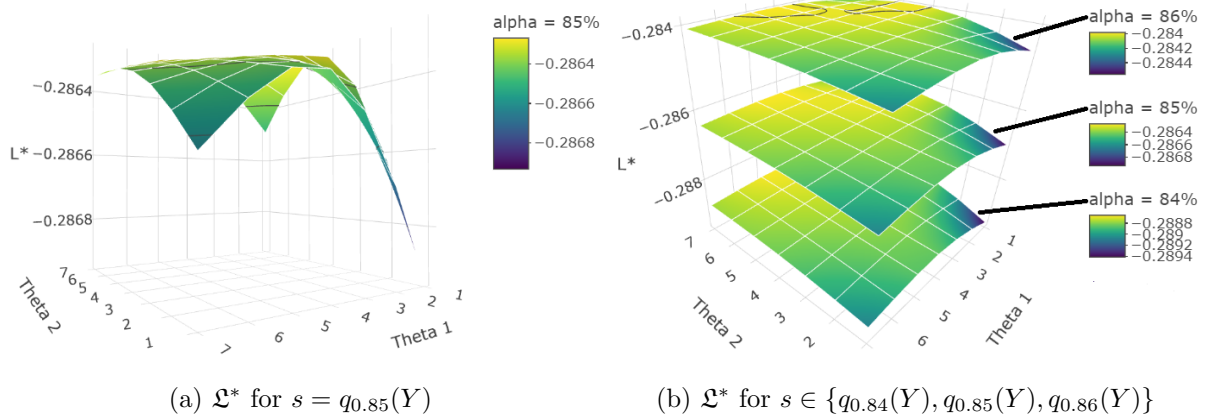


Figure 5: Objective function \mathcal{L}^* with respect to θ_1 and θ_2 for $s = q_{0.85}(Y)$ (panel (a)) and for $s \in \{q_{0.84}(Y), q_{0.85}(Y), q_{0.86}(Y)\}$ (panel (b)).

We conclude the analysis on real data by applying an algorithm similar to 1. Step 7 of this algorithm is slightly modified, as the three parameters a , b , and c of equation 5.1 are estimated in this case. The evolutions of these estimates with respect to sample size are shown in Figure 8 in the Appendix.

We observe a significant divergence in the parameter estimates for small sample sizes, underscoring the critical importance of data availability in constructing efficient index insurance models. However, the estimates stabilize and become acceptable at approximately half of the total sample size. This finding supports our proposed two-step approach to optimizing the objective function, even in scenarios with limited data.

For instance, in this example, estimating the distribution parameters using half of the available data in the first step would yield acceptable estimates of $\hat{S}(s|W)$ and $\hat{\gamma}(W)$. These estimates would then enable an efficient calculation of the objective function in the second step, facilitating the development of an index insurance model in a situation where half of the needed loss data Y is unavailable.

The results of applying this algorithm are presented in Figure 6. We observe a similar behavior to that seen with simulated data, confirming that the two-step estimation of \mathcal{L}^* described in steps 7 and 8 of the algorithm can improve the estimation of the objective function.

In summary, the proposed index insurance model for extreme losses offers a practical solution for covering such losses and introduces an extension that addresses the problem of economic loss

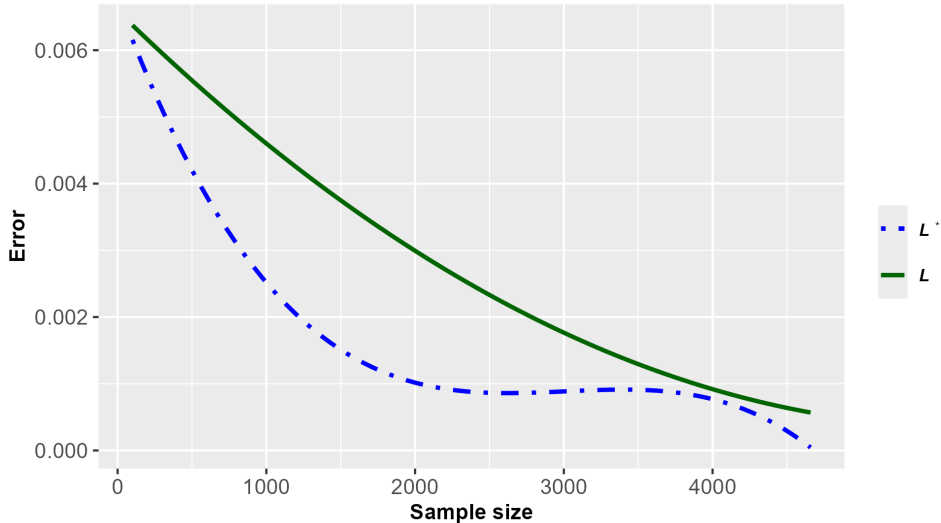


Figure 6: Errors in estimating \mathcal{L} by \mathcal{L}_n and \mathcal{L}_n^* as a function of sample size.

data scarcity—a major concern in insurance design. Indeed, in scenarios with scarce economic loss data Y , one could reasonably opt to work with $\hat{\mathcal{L}}_\theta^*$ even if the relative approximation error is higher than that of $\hat{\mathcal{L}}_\theta$, provided the error remains acceptable.

Data scarcity is particularly pronounced in developing countries, where data collection techniques and tools are often underdeveloped (Eze et al. [2020]). This lack of reliable data reduces the accuracy of insurance models in these regions, where certain insurance products could be vital given the significant risks faced. The methodology proposed in this paper addresses these challenges in two key ways:

- The lower cost of index insurance used in our model makes it more accessible and suitable for use in developing countries.
- The two-step estimation of the objective function \mathcal{L}^* provides a practical approach to mitigate the issue of data scarcity prevalent in these regions.

6 Comparison with a capped indemnity contract

In this section, we empirically investigate how the proposed hybrid contract positions itself compared to a capped indemnity contract. The capped indemnity contract (Zhou et al. [2010], Mao and Ostaszewski [2021]) is defined by the following payout:

$$X^{ST}(m) = Y\mathbf{1}_{Y \leq m} + m\mathbf{1}_{Y > m} = \min(Y, m).$$

Introducing a limit to compensation is a classical way to protect the insurer if the tail of the distribution is too heavy. On the other hand, this means that the policyholder will not benefit

from full coverage against the risk, and this also has consequences on the attractiveness of the insurance product.

On the other hand, the hybrid product with payout defined as

$$X^{HB}(s) = Y\mathbf{1}_{Y \leq s} + s\phi_{\theta}(\mathbf{W})\mathbf{1}_{Y > s},$$

is less limited in terms of maximum amount of compensation. The question is: for the same price, which type of contracts (X^{ST} or X^{HB}) will give the best coverage for the policyholder, in terms of average ratio between the compensation and the loss? For the same price, the threshold s of X^{HB} is necessarily smaller than the value of m used in X^{ST} , but this can be compensated by a better behavior of the parametric part.

In the rest of this section, we use the optimal values of the parameter θ obtained from the computations in the previous sections. The settings used in those sections for simulated and real losses are equally preserved here. To generalize this analysis, we assume that in the hybrid contract, the loading factors of the index part and that of the indemnity part of the contract are different. This is expected in practice due to the lower operational costs of index insurance (Barnett and Mahul [2007]). We denote by τ^i and τ^t the loading factors of index and traditional indemnity-based insurance, respectively.

The premium of the capped indemnity contract is therefore given by:

$$\pi^{SL}(m) = (1 + \tau^t)\mathbf{E}[X^{SL}] = (1 + \tau^t)\mathbf{E}[\min(Y, m)]$$

That of the hybrid contract is given by:

$$\pi^{HB}(s) = (1 + \tau^t)\mathbf{E}[Y\mathbf{1}_{Y \leq s}] + (1 + \tau^i)\mathbf{E}[s\phi_{\theta}(\mathbf{W})\mathbf{1}_{Y > s}]$$

We then define a function $m(s)$ such that, for a given s , $\pi^{HB}(s) = \pi^{SL}(m(s))$. The aim of this definition is to compare the expected rate of compensation of the capped indemnity contract ($\mathbf{E}[X^{SL}(m(s))/Y]$) to that of the hybrid contract ($\mathbf{E}[X^{HB}(s)/Y]$) at equal prices for a given s . For every s , the procedure consists in setting a value for τ^i , calculating the value of $\pi^{HB}(s)$, and using the function defined previously to find $m(s)$ such that $\pi^{HB}(s) = \pi^{SL}(m(s))$. Note that as τ^i varies, the rate of compensation of the hybrid contract remains unchanged (as s is fixed), but that of the capped indemnity contract has to adapt through $m(s)$ for the prices of the two contracts to be the same. The value of τ^t is kept fixed at 40%. Figure 7 shows the results obtained for simulated and real data.

Panels (a) and (c) of Figure 7 show that for large values of s (extreme losses) in the hybrid contract, the cap $m(s)$ of a capped indemnity contract of identical price is lower than s . Also, this value of $m(s)$ decreases as the loading of the index part of the proposed hybrid contract decreases. This highlights an additional advantage of the proposed hybrid contract in terms

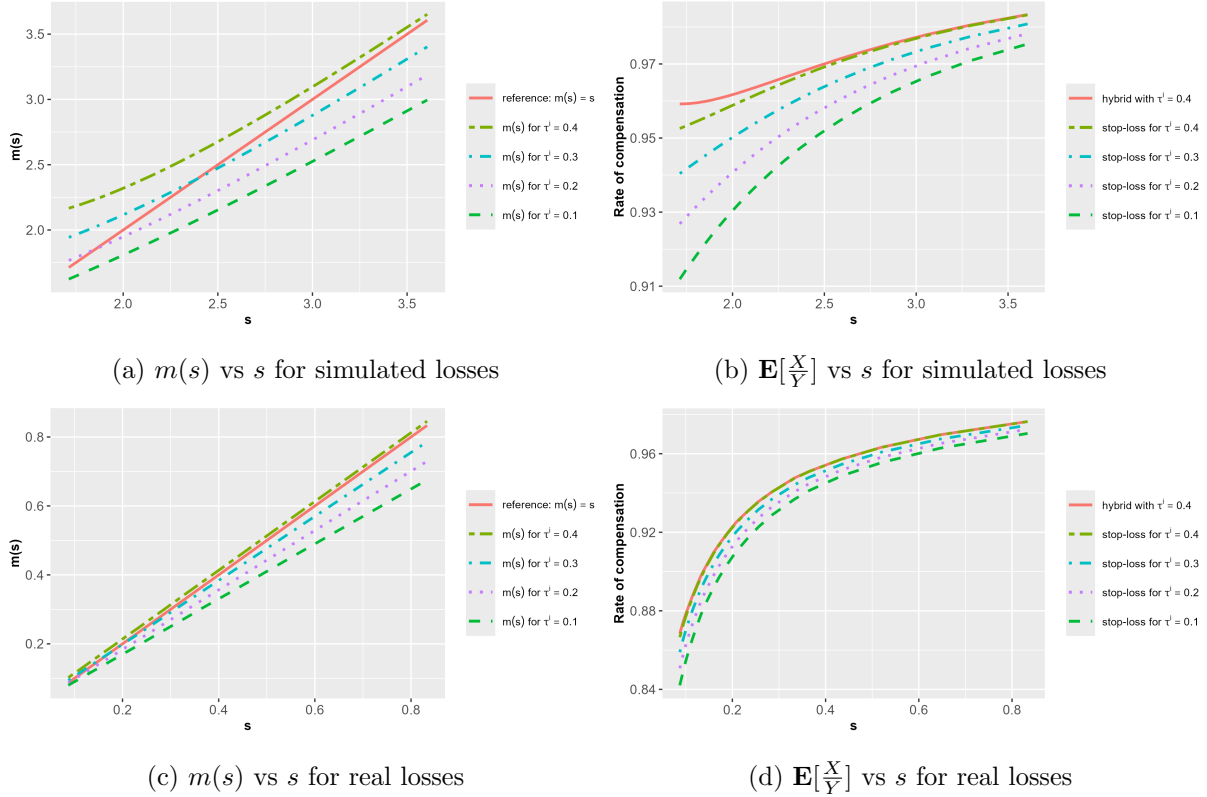


Figure 7: Comparative analysis of expected rates of compensation of the proposed hybrid insurance contract and the classical capped indemnity contract. The results are presented for real and simulated data and for various values of τ^i , the loading factor of index insurance. Note that τ^t is set at 40%.

of quality of coverage. Indeed, these results are presented in the form of expected rates of compensation in panels (b) and (d) of the same figure. These panels show that the expected rates of compensation of the capped indemnity contract decrease as the price of the index part of the hybrid contract decreases for a given value of s . This means that the advantage of the lower cost of index insurance makes the proposed hybrid contract more competitive and more suitable for policyholders than a capped indemnity contract of equal price. Moreover, further work on reducing the basis risk in the index part of this product could greatly improve this enhanced coverage provided by index insurance.

Also, recall that in the payout functions of the hybrid insurance contract presented in Sections 4.1 and 5.1, a condition is imposed such that compensation for losses above the threshold s is never less than s . This hybrid contract therefore has an implicit cap condition in addition to using index values to provide suitable compensation above s . This provides another explanation for why the hybrid contract compensates more than a classical capped indemnity contract of equal price.

7 Conclusion

In this paper, we studied the conception of an insurance product that mixes traditional insurance and an index-based part. The idea is to improve the ability to cover risks that are heavy tail, instead of applying a fixed limit of compensation. The use of covariates is used to provide an indemnity which is as close as possible to the needs of the policyholder, with a fast computation of the compensation which does not require expert analysis. The calibration of the index-based part of the product is based on a specific metric. Moreover, we provided an approximation of this criterion that allows to enhance the calibration part via the use of additional data on the covariates available after a claim (without the need of getting additional data on the economic losses).

The criterion that we use is of course dependent on assumptions on the behavior of the policyholders, and the question of the choice of the function L used for the optimization is a question that should be investigated from data on the behavior of the customers. This important question is beyond the scope of this paper. Regarding the behavior of the policyholder, we here consider that the only reason for preferring a reduced cover is a lower price. In practice, this may not be the sole factor: a faster compensation could be preferred, even though this compensation would be lower, since fast compensation may help the policyholder to repair faster, which is an important aspect of resilience. Moreover, even from a price prospective, the reduction of loading factors on the premiums of index insurance due to simplified claim management provides another level of optimization. This important effect of time is an aspect of the problem that could be investigated as an extension of the present work.

8 Appendix

The Appendix section is organized as follows. In section 9, we study the comparison between the product with "ideal" trigger mechanism and the one where the trigger is based only on the covariates (Proposition 2.1). In section 9.1, we provide the proof of the key approximation result, that is Theorem 9.1. Section 9.2 is related to the proof of Theorem 3.2 on the convergence rate of the index optimization procedure, which relies on a technical Lemma presented in section 9.3. Additional empirical results related to the real data analysis are provided in section 9.4.

9 Proof of Proposition 2.1

$$X_{\tilde{\theta}(s)}(s) - X^*(s, \tilde{s}) = (Y - s\phi_{\tilde{\theta}(s)}(\mathbf{W}))\mathbf{1}_{E_-(s, \tilde{s})} + (s\phi_{\tilde{\theta}(s)}(\mathbf{W}) - Y)\mathbf{1}_{E_+(s, \tilde{s})},$$

where

$$\begin{aligned} E_-(s, \tilde{s}) &= \left\{ Y \leq s \text{ and } s\phi_{\tilde{\theta}(s)} \geq \tilde{s} \right\}, \\ E_+(s, \tilde{s}) &= \left\{ Y > s \text{ and } s\phi_{\tilde{\theta}(s)} \leq \tilde{s} \right\}. \end{aligned}$$

Then,

$$X_{\tilde{\theta}(s)}(s) - X^*(s, \tilde{s}) \leq (s - \tilde{s})\mathbf{1}_{E_-(s, \tilde{s})}.$$

Then,

$$\begin{aligned} E \left[L \left(\frac{X^*(s, \tilde{s})}{Y} - f(\pi^*(s, \tilde{s})) \right) \right] &\geq E \left[L \left(\frac{X_{\tilde{\theta}(s)}(s)}{Y} - f(\pi^*(s)) \right) \mathbf{1}_{\bar{E}_-(s, \tilde{s})} \right] \\ &\quad + E \left[L \left(\frac{(s - \tilde{s})}{Y} - f(\pi^*(s)) \right) \mathbf{1}_{E_-(s, \tilde{s})} \right], \end{aligned}$$

where \bar{A} denotes the complementary of the set A . This leads to

$$\begin{aligned} E \left[L \left(\frac{X^*(s, \tilde{s})}{Y} - f(\pi^*(s)) \right) \right] &\geq E \left[L \left(\frac{X_{\tilde{\theta}(s)}(s)}{Y} - f(\pi^*(s, \tilde{s})) \right) \right] \\ &\quad - E \left[\left\{ L(1 - f(\pi^*(s, \tilde{s}))) - L \left(\frac{(s - \tilde{s})}{Y} - f(\pi^*(s, \tilde{s})) \right) \right\} \mathbf{1}_{E_-(s, \tilde{s})} \right] \\ &\geq E \left[L \left(\frac{X_{\tilde{\theta}(s)}(s)}{Y} - f(\pi^*(s, \tilde{s})) \right) \right] \\ &\quad - \mathbb{P}(E_-(s, \tilde{s})) [L(1 - f(\pi^*(s, \tilde{s}))) - L(-f(\pi^*(s, \tilde{s})))]. \end{aligned}$$

Moreover, from a Taylor expansion,

$$E \left[L \left(\frac{X_{\tilde{\theta}(s)}(s)}{Y} - f(\pi^*(s, \tilde{s})) \right) \right] \geq E \left[L \left(\frac{X_{\tilde{\theta}(s)}(s)}{Y} - f(\pi_{\tilde{\theta}}(s)) \right) \right] - \|L'\|_\infty \|f'\|_\infty \times (\pi^*(s, \tilde{s}) - \pi_{\tilde{\theta}}(s))_+.$$

We have

$$\pi^*(s, \tilde{s}) - \pi_{\tilde{\theta}}(s) \leq E \left[(s\phi_{\tilde{\theta}(s)}(\mathbf{W}) - Y)_+ \mathbf{1}_{E_-(s, \tilde{s})} \right],$$

and the result follows.

9.1 Proof of Theorem 3.1

Let

$$\mathfrak{L}_\theta(s|\mathbf{w}) = E \left[L \left(\frac{X_\theta(s)}{Y} - f(\pi_\theta(s)) \right) \mid \mathbf{W} = \mathbf{w} \right].$$

We have

$$\begin{aligned}
\mathcal{L}_\theta(s|\mathbf{w}) &= E \left[L \left(\mathbf{1}_{Y \leq s} + \frac{s\phi(\theta, \mathbf{w})}{Y} \mathbf{1}_{Y > s} - f(\pi_\theta(s)) \right) \mid \mathbf{W} = \mathbf{w} \right] \\
&= \mathbb{P}(Y \leq s \mid \mathbf{W} = \mathbf{w}) L(1 - f(\pi_\theta(s))) - \int_s^\infty L \left(\frac{s\phi(\theta, \mathbf{w})}{t} - f(\pi_\theta(s)) \right) dS_Y(t|\mathbf{w}) \\
&= L(1 - f(\pi_\theta(s))) \left\{ 1 - S(s|\mathbf{w}) \left[1 - \frac{L(\phi(\theta, \mathbf{w}) - f(\pi_\theta(s)))}{L(1 - f(\pi_\theta(s)))} + \mathcal{L}_{\theta,1}(s|\mathbf{w}) \right] \right\}, \quad (9.1)
\end{aligned}$$

with

$$\begin{aligned}
\mathcal{L}_{\theta,1}(s|\mathbf{w}) &= s\phi(\theta, \mathbf{w}) \int_s^\infty \frac{L' \left(\frac{s\phi(\theta, \mathbf{w})}{t} - f(\pi_\theta(s)) \right)}{L(1 - f(\pi_\theta(s)))} \frac{S(t|\mathbf{w}) dt}{S(s|\mathbf{w}) t^2} \\
&= \phi(\theta, \mathbf{w}) \int_1^\infty \frac{L' \left(\frac{\phi(\theta, \mathbf{w})}{u} - f(\pi_\theta(s)) \right)}{L(1 - f(\pi_\theta(s)))} \frac{l(su|\mathbf{w})}{l(s|\mathbf{w})} \frac{du}{u^{2+\frac{1}{\gamma(\mathbf{w})}}}.
\end{aligned}$$

From Assumption 2,

$$\frac{L(\phi(\theta, \mathbf{w}) - f(\pi_\theta(s)))}{L(1 - f(\pi_\theta(s)))} = \varphi_0(\phi(\theta, \mathbf{w}))(1 + R_{\theta,0}(s, \mathbf{w})) \quad (9.2)$$

with $\sup_{\theta, \mathbf{w}} |R_{\theta,0}(s, \mathbf{w})| \rightarrow 0$, when s tends to infinity.

We have,

$$\mathcal{L}_{\theta,1}(s|\mathbf{w}) = \phi(\theta, \mathbf{w}) \int_1^\infty \varphi_1(\phi(\theta, \mathbf{w})/u) \frac{l(su|\mathbf{w}) du}{l(s|\mathbf{w}) u^{2+\frac{1}{\gamma(\mathbf{w})}}} (1 + R_{\theta,1}(s, \mathbf{w})),$$

where, from Assumption 1,

$$|R_{\theta,1}(s, \mathbf{w})| \leq \frac{C}{c} \left\{ \sup_t \frac{\left| \frac{L'(t - f(\pi_\theta(s)))}{L(1 - f(\pi_\theta(s)))} - \varphi_1(t) \right|}{\varphi_1(t)} \right\} \int_1^\infty \frac{l(su)}{l(s)} \frac{du}{u^2}.$$

Since l is slow-varying, dominated convergence yields

$$\int_1^\infty \frac{l(su)}{l(s)} \frac{du}{u^2} \rightarrow \int_1^\infty \frac{du}{u^2},$$

when s tends to infinity. Then, from Assumption 3, we get

$$\sup_{\theta, \mathbf{w}} |R_{\theta,1}(s, \mathbf{w})| = o(1). \quad (9.3)$$

Next, let

$$\mathcal{L}_{\theta,2}(s|\mathbf{w}) = \phi(\theta, \mathbf{w}) \int_1^\infty \varphi_1(\phi(\theta, \mathbf{w})/u) \frac{l(su|\mathbf{w}) du}{l(s|\mathbf{w}) u^{2+\frac{1}{\gamma(\mathbf{w})}}}.$$

We have

$$\begin{aligned}
\mathcal{L}_{\theta,2}(s|\mathbf{w}) &= \phi(\theta, \mathbf{w}) \int_1^\infty \varphi_1(\phi(\theta, \mathbf{w})/u) \frac{du}{u^{2+\frac{1}{\gamma(\mathbf{w})}}} + R_{\theta,2}(s, \mathbf{w}) \\
&= \left\{ \phi(\theta, \mathbf{w})^{-1/\gamma(\mathbf{w})} \int_0^{\phi(\theta, \mathbf{w})} v^{1/\gamma(\mathbf{w})} \varphi_1(v) dv \right\} (1 + R_{\theta,2}(s, \mathbf{w})) \\
&= \Phi_1(\phi(\theta, \mathbf{w}), \gamma(\mathbf{w})),
\end{aligned}$$

with

$$R_{\theta,2}(s, \mathbf{w}) = \phi(\theta, \mathbf{w}) \int_1^\infty \varphi_1(\phi(\theta, \mathbf{w})/u) \left[\frac{l(su|\mathbf{w})}{l(s|\mathbf{w})} - 1 \right] \frac{du}{u^{2+\frac{1}{\gamma(\mathbf{w})}}}.$$

We have

$$R_{\theta,2}(s, \mathbf{w}) \leq |\Phi_1(\phi(\theta, \mathbf{w}), \gamma(\mathbf{w}))| \sup_{\mathbf{w}} \left| \frac{l(su|\mathbf{w})}{l(s|\mathbf{w})} - 1 \right|,$$

and, from Assumption 1,

$$\sup_{\theta, \mathbf{w}} \left| \frac{R_{\theta,2}(s, \mathbf{w})}{\Phi_1(\phi(\theta, \mathbf{w}), \gamma(\mathbf{w}))} \right| \xrightarrow{s \rightarrow \infty} 0. \quad (9.4)$$

Gathering (9.1) to (9.4), we get

$$\mathfrak{L}_\theta(s|\mathbf{w}) = L(1 - f(\pi_\theta(s))) \{1 - S(s|\mathbf{w}) [1 - \varphi_0(\phi(\theta, \mathbf{w})) + \Phi_1(\phi(\theta, \mathbf{w}))]\} (1 + R_\theta(s, \mathbf{w})),$$

with $\sup_{\theta, w} |R(s, \mathbf{w})| \rightarrow 0$ when s tends to infinity.

The result then follows, since $\mathfrak{L}_\theta(s) = E[\mathfrak{L}_\theta(s)|\mathbf{W}]$.

9.2 Proof of Theorem 3.2

Let us denote

$$\bar{\mathfrak{L}}_\theta^*(s) = \frac{1}{m} \sum_{j=1}^m \Psi(\pi_\theta(s), \phi(\theta, \mathbf{W}_j); S(s|\mathbf{W}_j), \gamma(\mathbf{W}_j)).$$

This function of θ is the one we could compute if we knew exactly S and γ . The proof Theorem 3.2 is essentially based on the fact that the difference between $\bar{\mathfrak{L}}_\theta^*(s)$ and $\hat{\mathfrak{L}}_\theta^*(s)$ is asymptotically negligible. Thus, $\bar{\theta}^*(s) = \arg \max_{\theta \in \Theta} \bar{\mathfrak{L}}_\theta^*(s)$ and $\hat{\theta}^*(s)$ are asymptotically equivalent.

The result of the Theorem is shown in two steps. First, we show the consistency of $\hat{\theta}^*(s)$ (Proposition 9.1), then the convergence rate is derived in a second step.

Step 1: consistency.

Proposition 9.1 *Under the assumptions of Theorem 3.2,*

$$\hat{\theta}^*(s) - \bar{\theta}^*(s) = o_P(1).$$

Proof. To prove the consistency, it suffices to show that

$$\sup_{\theta} |\hat{\mathfrak{L}}_\theta^*(s) - \mathfrak{L}_\theta^*(s)| = o_P(1),$$

since $\tilde{\theta}^*(s)$ is the unique maximizer of $\mathfrak{L}_\theta^*(s)$ and is in the interior part of Θ (see for example Theorem 5.7 in Van der Vaart [2000]).

We have

$$\sup_{\theta \in \Theta} \left| \hat{\mathfrak{L}}_\theta^*(s) - \mathfrak{L}_\theta^*(s) \right| \leq \sup_{\theta \in \Theta} \left| \hat{\mathfrak{L}}_\theta^*(s) - \bar{\mathfrak{L}}_\theta^*(s) \right| + \sup_{\theta \in \Theta} \left| \bar{\mathfrak{L}}_\theta^*(s) - \mathfrak{L}_\theta^*(s) \right|. \quad (9.5)$$

We have

$$\left| \Psi(\pi_\theta(s), \phi(\theta, \mathbf{w}); S(s|\mathbf{w}), \gamma(\mathbf{w})) - \Psi(\pi_{\theta'}(s), \phi(\theta', \mathbf{w}); S(s|\mathbf{w}), \gamma(\mathbf{w})) \right| \leq \|\theta - \theta'\| L(\mathbf{w}),$$

with

$$\begin{aligned} L(\mathbf{w}) &= \sup_{\theta \in \Theta} \left| -f'(\pi_\theta(s)) \frac{\partial \pi_\theta(s)}{\partial \theta} L'(1 - f(\pi_\theta(s))) (1 - S(s|\mathbf{w})) \Phi_0(\phi(\theta, \mathbf{w}), \gamma(\mathbf{w})) \right| \\ &\quad + \sup_{\theta \in \Theta} \left| \frac{\partial \phi(\theta, \mathbf{w})}{\partial \theta} L(1 - f(\pi_\theta(s))) (1 - S(s|\mathbf{w})) \partial_1 \Phi_0(\phi(\theta, \mathbf{w}), \gamma(\mathbf{w})) \right| \\ &\leq C \Lambda(\mathbf{w}), \end{aligned}$$

from Assumption 4, for some constant C .

Therefore, from Example 19.7 in Van der Vaart [2000], we get

$$\sup_{\theta \in \Theta} \left| \bar{\mathfrak{L}}_\theta^*(s) - \mathfrak{L}_\theta^*(s) \right| = o_P(1). \quad (9.6)$$

Next,

$$\begin{aligned} \hat{\mathfrak{L}}_\theta^*(s) - \bar{\mathfrak{L}}_\theta^*(s) &= \frac{1}{m} \sum_{j=1}^m \left\{ \frac{S(s|\mathbf{W}_j) - \hat{S}(s|\mathbf{W}_j)}{1 - S(s|\mathbf{W}_j)} + \int_0^{\phi(\theta, \mathbf{W}_j)} \left(v^{1/\hat{\gamma}(\mathbf{W}_j)} - v^{1/\gamma(\mathbf{W}_j)} \right) \varphi_1(v) dv \right\} \\ &\quad \times \Psi(\pi_\theta(s), \phi(\theta, \mathbf{W}_j); S(s|\mathbf{W}_j), \gamma(\mathbf{W}_j)). \end{aligned}$$

We get, for some constant C_0 ,

$$\left| \hat{\mathfrak{L}}_\theta^*(s) - \bar{\mathfrak{L}}_\theta^*(s) \right| \leq C_0 \sup_{\mathbf{w} \in \mathcal{W}} \left| \frac{S(s|\mathbf{w}) - \hat{S}(s|\mathbf{w})}{1 - S(s|\mathbf{w})} \right| \times \sup_{\mathbf{w} \in \mathcal{W}} \left| \frac{1}{\hat{\gamma}(\mathbf{w})} - \frac{1}{\gamma(\mathbf{w})} \right| \int_0^\infty v^{1/\gamma} \varphi_1(v) dv,$$

since Ψ is bounded. Assumption 5 and Assumption 4 lead to

$$\sup_{\theta \in \Theta} \left| \hat{\mathfrak{L}}_\theta^*(s) - \bar{\mathfrak{L}}_\theta^*(s) \right| = o_P(1). \quad (9.7)$$

Merging (9.5) to (9.7), we get the result. ■

Step 2: Rate of consistency.

To shorten the notation, $\hat{\theta}^*$ (resp. $\bar{\theta}^*$, $\tilde{\theta}^*$) will denote $\hat{\theta}^*(s)$ (resp. $\bar{\theta}^*(s)$, $\tilde{\theta}^*(s)$) in this section. Let ∇_θ denote the gradient vector of a function with respect to θ . By definition,

$$\nabla_\theta \hat{\mathfrak{L}}_{\hat{\theta}^*}^*(s) = 0.$$

From a Taylor expansion and the consistency of $\hat{\theta}^*$,

$$0 = \nabla_{\theta} \hat{\mathcal{L}}_{\hat{\theta}^*}^*(s) + [\tilde{\theta}^* - \hat{\theta}^*][\nabla_{\theta}^2 \hat{\mathcal{L}}_{\hat{\theta}^*}^*(s) + o_P(1)]. \quad (9.8)$$

From the consistency of \hat{S} and $\hat{\gamma}$ in Assumption 5 and the law of large numbers, we get

$$\nabla_{\theta}^2 \hat{\mathcal{L}}_{\hat{\theta}^*}^*(s) = \nabla_{\theta}^2 \bar{\mathcal{L}}_{\hat{\theta}^*}^*(s) + o_P(1) = \nabla_{\theta}^2 \mathcal{L}_{\hat{\theta}^*}^*(s) + o_P(1) = \Sigma(s) + o_P(1).$$

Next,

$$\begin{aligned} \nabla_{\theta} \hat{\mathcal{L}}_{\hat{\theta}^*}^*(s) &= \nabla_{\theta} \bar{\mathcal{L}}_{\hat{\theta}^*}^*(s) + \frac{1}{m} \sum_{j=1}^m \left\{ \nabla_{\theta} \Psi \left(\pi_{\tilde{\theta}^*}(s), \phi(\tilde{\theta}^*, \mathbf{W}_j); \hat{S}(s|\mathbf{W}_j), \hat{\gamma}(\mathbf{W}_j) \right) \right. \\ &\quad \left. - \nabla_{\theta} \Psi \left(\pi_{\tilde{\theta}^*}(s), \phi(\tilde{\theta}^*, \mathbf{W}_j); S(s|\mathbf{W}_j), \gamma(\mathbf{W}_j) \right) \right\}. \end{aligned}$$

From Assumption 6 and Lemma 9.2, we get, from the asymptotic equicontinuity property of Donsker classes (see section 2.8.2 in Van Der Vaart et al. [1996]),

$$\begin{aligned} &\frac{1}{m} \sum_{j=1}^m \left\{ \nabla_{\theta} \Psi \left(\pi_{\tilde{\theta}^*}(s), \phi(\tilde{\theta}^*, \mathbf{W}_j); \hat{S}(s|\mathbf{W}_j), \hat{\gamma}(\mathbf{W}_j) \right) \right. \\ &\quad \left. - \nabla_{\theta} \Psi \left(\pi_{\tilde{\theta}^*}(s), \phi(\tilde{\theta}^*, \mathbf{W}_j); S(s|\mathbf{W}_j), \gamma(\mathbf{W}_j) \right) \right\} \\ &= \int \left\{ \nabla_{\theta} \Psi \left(\pi_{\tilde{\theta}^*}(s), \phi(\tilde{\theta}^*, \mathbf{w}); \hat{S}(s|\mathbf{w}), \hat{\gamma}(\mathbf{w}) \right) \right\} d\mathbb{P}(\mathbf{w}) + o_P(m^{-1/2}) = B_n + o_P(m^{-1/2}). \end{aligned}$$

Back to (9.8), we get

$$\tilde{\theta}^* - \hat{\theta}^* = \Sigma^{-1} \left\{ \frac{1}{m} \sum_{j=1}^m \nabla_{\theta} \Psi \left(\pi_{\tilde{\theta}^*}(s), \phi(\tilde{\theta}^*, \mathbf{W}_j); S(s|\mathbf{W}_j), \gamma(\mathbf{W}_j) \right) \right\} + O_P(B_n) + o_P(m^{-1/2}).$$

9.3 Technical lemma on Donsker classes

Lemma 9.2 *Let*

$$\mathcal{A} = \{ \mathbf{w} \rightarrow \nabla_{\theta} \Psi(\pi_{\theta}(s), \phi(\theta, \mathbf{w}); h_1(s, \mathbf{w}), h_2(\mathbf{w})) : h_1 \in \mathcal{H}_1, h_2 \in \mathcal{H}_2, \theta \in \Theta \}.$$

Then \mathcal{A} is Donsker under Assumption 6.

Proof. Let

$$\mathcal{A}_3 = \left\{ \mathbf{w} \rightarrow \frac{\partial \phi(\theta, \mathbf{w})}{\partial \theta} : \theta \in \Theta \right\}.$$

We have

$$\begin{aligned} &\nabla_{\theta} \Psi(\pi_{\theta}(s), \phi(\theta, \mathbf{w}); h_1(s, \mathbf{w}), h_2(\mathbf{w})) = \\ &-f'(\pi_{\theta}(s)) \frac{\partial \pi_{\theta}(s)}{\partial \theta} L'(1 - f(\pi_{\theta}(s)))(1 - h_1(s, \mathbf{w})) \Phi_0(\phi(\theta, \mathbf{w}), h_2(\mathbf{w})) \\ &\quad + \frac{\partial \phi(\theta, \mathbf{w})}{\partial \theta} L(1 - f(\pi_{\theta}(s)))(1 - h_1(s, \mathbf{w})) \partial_1 \Phi_0(\phi(\theta, \mathbf{w}), h_2(\mathbf{w})). \end{aligned}$$

The classes of functions $\mathcal{A}_1 = \{\mathbf{w} \rightarrow \Phi_0(\phi(\theta, \mathbf{w}), h_2(\mathbf{w})) : h_2 \in \mathcal{H}_2\}$ and $\mathcal{A}_2 = \{\mathbf{w} \rightarrow \partial_1 \Phi_0(\phi(\theta, \mathbf{w}), h_2(\mathbf{w})) : h_2 \in \mathcal{H}_2\}$ inherit the Donsker property of classes $\{\mathbf{w} \rightarrow \phi(\theta, \mathbf{w}) : \theta \in \Theta\}$ and \mathcal{H}_2 from Theorem 2.10.6 in Van Der Vaart et al. [1996]. Indeed,

$$\begin{aligned} |\Phi_0(\phi(\theta, \mathbf{w}), h_2(\mathbf{w})) - \Phi_0(\phi(\theta', \mathbf{w}), h_2'(\mathbf{w}))| &\leq |\varphi_0(\phi(\theta, \mathbf{w})) - \varphi_0(\phi(\theta', \mathbf{w}))| + \left| \int_{\phi(\theta, \mathbf{w})}^{\phi(\theta', \mathbf{w})} v^{1/\gamma-} \varphi_1(v) dv \right| \\ &\quad + \left| \frac{1}{h_2(\mathbf{w})} - \frac{1}{h_2'(\mathbf{w})} \right| \int_0^\infty v^\alpha \varphi_1(v) dv \\ &\leq C (|\phi(\theta, \mathbf{w}) - \phi(\theta', \mathbf{w})| + |h_2(\mathbf{w}) - h_2'(\mathbf{w})|), \end{aligned}$$

for some constant C , and

$$\begin{aligned} &|\partial_1 \Phi_0(\phi(\theta, \mathbf{w}), h_2(\mathbf{w})) - \partial_1 \Phi_0(\phi(\theta', \mathbf{w}), h_2'(\mathbf{w}))| \leq \\ &|\varphi_0(\phi(\theta, \mathbf{w})) - \varphi_0(\phi(\theta', \mathbf{w}))| \\ &+ \left| \phi(\theta, \mathbf{w})^{1/h_2(\mathbf{w})} \varphi_1(\phi(\theta, \mathbf{w})) - \phi(\theta, \mathbf{w})^{1/h_2'(\mathbf{w})} \varphi_1(\phi(\theta', \mathbf{w})) \right| \\ &\leq C' (|\phi(\theta, \mathbf{w}) - \phi(\theta', \mathbf{w})| + |h_2(\mathbf{w}) - h_2'(\mathbf{w})|), \end{aligned}$$

for a constant C' . Then, the classes $(1 - \mathcal{H}_1)\mathcal{A}_1$ and $\mathcal{A}_3(1 - \mathcal{H}_1)\mathcal{A}_2$ are Donsker from the boundedness of \mathcal{H}_1 and \mathcal{A}_3 and Example 2.10.8 in Van Der Vaart et al. [1996]. Note that the boundedness of \mathcal{A}_3 can be weakened by additional conditions on \mathcal{A}_2 : the key issue is that the class $\mathcal{A}_3(1 - \mathcal{H}_1)\mathcal{A}_2$ should be dominated by a L^2 random variable. Then, the class \mathcal{A} is included in the addition of these two classes (Donsker from Example 2.10.7 in Van Der Vaart et al. [1996]), up to multiplications by uniformly bounded quantities depending on θ (but not on \mathbf{w}). Theorem 2.10.6 in Van Der Vaart et al. [1996] applies again, leading to the result. ■

9.4 Additional results on real data

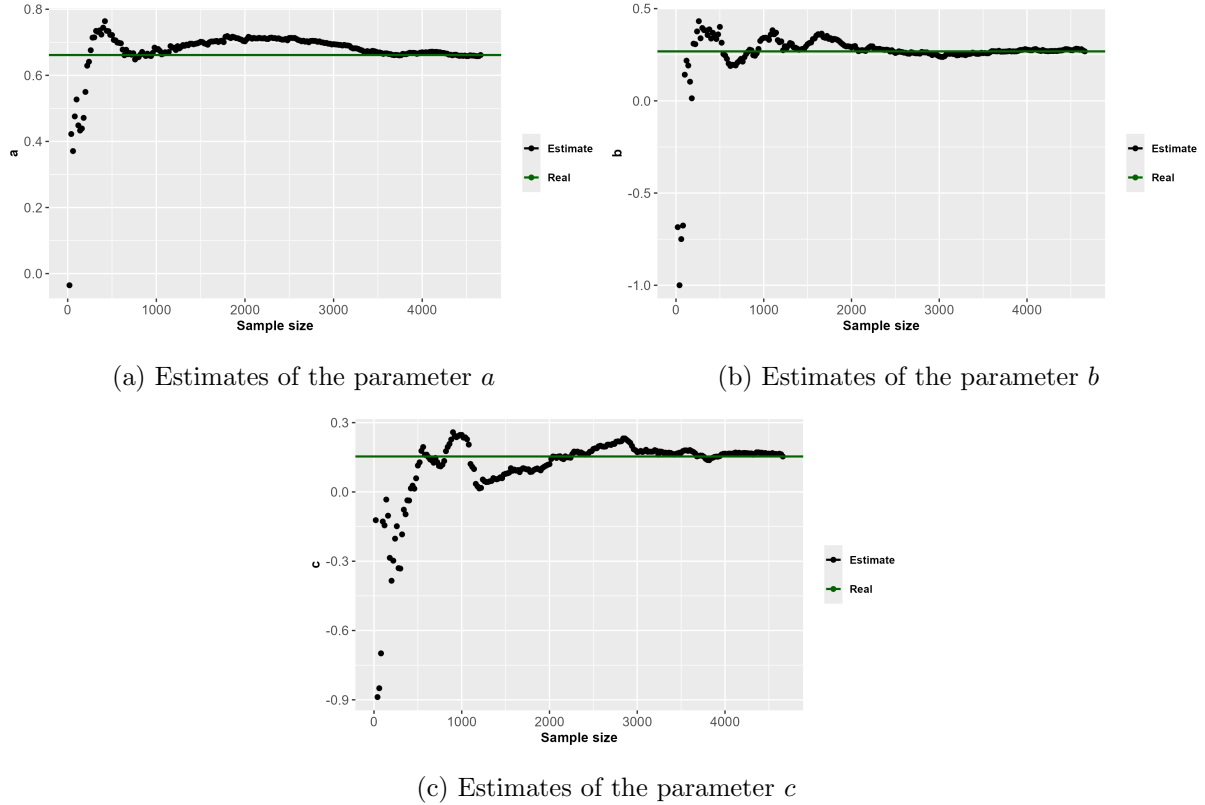


Figure 8: Evolutions of the estimates of the parameters a (panel (a)), b (panel (b)) and c (panel(c)) of equation 5.1 with sample size.

Acknowledgment: Olivier Lopez acknowledges funding from the Excellence Chair CARE (Allianz, Ensaie, Risk Fundation).

References

- M. J. Abdi, N. Raffar, Z. Zulkaffli, K. Nurulhuda, B. M. Rehan, F. M. Muharam, N. A. Khosim, and F. Tangang. Index-based insurance and hydroclimatic risk management in agriculture: A systematic review of index selection and yield-index modelling methods. *International Journal of Disaster Risk Reduction*, 67:102653, 2022.
- H. Alderman and T. Haque. *Insurance against covariate shocks: The role of index-based insurance in social protection in low-income countries of Africa*. Number 95. World Bank Publications, 2007.
- C. André, D. Monfort, M. Bouzit, and C. Vinchon. Contribution of insurance data to cost assessment of coastal flood damage to residential buildings: insights gained from johanna

- (2008) and xynthia (2010) storm events. *Natural Hazards and Earth System Sciences*, 13(8): 2003–2012, 2013.
- B. J. Barnett and O. Mahul. Weather index insurance for agriculture and rural areas in lower-income countries. *American Journal of Agricultural Economics*, 89(5):1241–1247, 2007.
- J. Beirlant, Y. Goegebeur, J. Segers, and J. L. Teugels. *Statistics of extremes: theory and applications*, volume 558. John Wiley & Sons, 2004.
- A. Braun, M. Eling, and C. Jaenicke. Cyber insurance-linked securities. *ASTIN Bulletin: The Journal of the IAA*, 53(3):684–705, 2023.
- M. Carter, A. de Janvry, E. Sadoulet, and A. Sarris. Index insurance for developing country agriculture: a reassessment. *Annual Review of Resource Economics*, 9(1):421–438, 2017.
- Z. Chen, Y. Lu, J. Zhang, and W. Zhu. Managing weather risk with a neural network-based index insurance. *Management Science*, 70(7):4306–4327, 2024.
- S. Conradt, R. Finger, and R. Bokusheva. Tailored to the extremes: Quantile regression for index-based insurance contract design. *Agricultural economics*, 46(4):537–547, 2015.
- A. Daouia, G. Stupfler, and A. Usseglio-Carleve. Inference for extremal regression with dependent heavy-tailed data. *The Annals of Statistics*, 51(5):2040–2066, 2023.
- N. A. Doherty and A. Richter. Moral hazard, basis risk, and gap insurance. *Journal of Risk and Insurance*, 69(1):9–24, 2002.
- G. Elabed, M. F. Bellemare, M. R. Carter, and C. Guirkinger. Managing basis risk with multi-scale index insurance. *Agricultural Economics*, 44(4-5):419–431, 2013.
- E. Eze, A. Girma, A. Zenebe, J. M. Kourouma, and G. Zenebe. Exploring the possibilities of remote yield estimation using crop water requirements for area yield index insurance in a data-scarce dryland. *Journal of Arid Environments*, 183:104261, 2020.
- S. Farkas, O. Lopez, and M. Thomas. Cyber claim analysis using generalized pareto regression trees with applications to insurance. *Insurance: Mathematics and Economics*, 98:92–105, 2021.
- S. Farkas, A. Heranval, O. Lopez, and M. Thomas. Generalized pareto regression trees for extreme event analysis. *Extremes*, pages 1–41, 2024.
- Z. Gu, Y. Li, M. Zhang, and Y. Liu. Modelling economic losses from earthquakes using regression forests: Application to parametric insurance. *Economic Modelling*, 125:106350, 2023. ISSN

0264-9993. doi: <https://doi.org/10.1016/j.econmod.2023.106350>. URL <https://www.sciencedirect.com/science/article/pii/S0264999323001621>.

- W. K. Härdle and B. L. Cabrera. Calibrating cat bonds for mexican earthquakes. *Journal of Risk and Insurance*, 77(3):625–650, 2010.
- X. Lin and W. J. Kwon. Application of parametric insurance in principle-compliant and innovative ways. *Risk Management and Insurance Review*, 23(2):121–150, 2020.
- A. Louaas and P. Picard. Optimal insurance coverage of low-probability catastrophic risks. *The Geneva Risk and Insurance Review*, 46:61–88, 2021.
- H. Mao and K. Ostaszewski. Optimal claim settlement strategies under constraint of cap on claim loss. *Mathematics*, 9(24):3284, 2021.
- L. Mastroeni, A. Mazzocchi, and M. Naldi. Pricing cat bonds for cloud service failures. *Journal of Risk and Financial Management*, 15(10):463, 2022.
- E. Michel-Kerjan and F. Morlaye. Extreme events, global warming, and insurance-linked securities: How to trigger the “tipping point”. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 33:153–176, 2008.
- I.-L. Securities. The handbook of insurance-linked securities. 2009.
- K. S. Tan and J. Zhang. Flexible weather index insurance design with penalized splines. *North American Actuarial Journal*, 28(1):1–26, 2024.
- A. W. Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.
- A. W. Van Der Vaart, J. A. Wellner, A. W. van der Vaart, and J. A. Wellner. *Weak convergence and Empirical Processes*. Springer, 1996.
- C. Zhou, W. Wu, and C. Wu. Optimal insurance in the presence of insurer’s loss limit. *Insurance: Mathematics and Economics*, 46(2):300–307, 2010.