



HAL
open science

Index insurance under demand and solvency constraints

Olivier Lopez, Daniel Nkameni

► **To cite this version:**

Olivier Lopez, Daniel Nkameni. Index insurance under demand and solvency constraints. 2025. ⟨hal-05182769v2⟩

HAL Id: hal-05182769

<https://hal.science/hal-05182769v2>

Preprint submitted on 23 Oct 2025 (v2), last revised 16 Feb 2026 (v3)

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

Index insurance under demand and solvency constraints

Olivier LOPEZ¹, Daniel NKAMENI^{1,2}

October 23, 2025

Abstract

Index insurance is often proposed to reduce protection gaps, especially for emerging risks. Unlike traditional insurance, it bases compensation on a measurable index, enabling faster payouts and lower claim management costs. This approach benefits both policyholders, through quick payments, and insurers, through reduced costs and better risk control due to reliable data and robust statistical estimates. An important difference with the concept of Cat Bonds is that the feasibility of such coverage relies on the possibility of mutualization. Mutualization, in turn, is achieved only if a sufficiently high number of policyholders agree to subscribe. The purpose of this paper is to introduce a model for the demand for index insurance and to provide conditions under which the solvency of the portfolio is achieved. From these conditions, we deduce a product that combines index and traditional indemnity insurance in order to benefit from the best of both approaches. We illustrate our results with a practical example involving the design of an index insurance product in the field of cyber insurance.

Key words: Index insurance; expected utility; demand in insurance; cyber insurance

JEL code: G220

Short title: Demand for index insurance.

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

Index insurance is often promoted as a solution capable of addressing certain structural weaknesses of traditional indemnity-based insurance (see for example Barnett and Mahul [2007], Carter et al. [2017], Prokopchuk et al. [2018] or Han et al. [2019]). The principle of these coverages lies in the use of an index (or parameter) that can be easily calculated based on information available immediately after the incident. This automated calculation significantly simplifies claim management. Since the compensation is determined without the need for an expert to evaluate the amount, it can be paid out very quickly to the policyholder. The clarity of the indemnification conditions also reduces the likelihood of legal disputes. On the other hand, the policyholder must bear basis risk (see Clement et al. [2018]), as the compensation is not based on the actual loss but instead on an approximation of the latter. Consequently, there is a concern that index insurance may have a disappointing aspect, which is seen as a barrier to its development (see Johnson [2021]).

The design of an index that can serve as the basis for insurance coverage is similar to a statistical problem of estimation or prediction: using available variables, the goal is to approximate as closely as possible an unobserved quantity (the loss experienced by the policyholder), see Cesarini et al. [2021b]. However, this problem involves a number of constraints due to the need to align with the policyholder's expectations. For instance, the construction of an index proposed by Conradt et al. [2015], Zhang et al. [2019], or Chen et al. [2023] is based on maximizing the policyholder's utility rather than relying on a more standard metric commonly used in regression or forecasting.

Indeed, behind the question of meeting the policyholder's needs lies the issue of demand. Index insurance remains an insurance product: while other index-based products, such as Cat Bonds, can achieve balance through diversification strategies inherent to financial instruments, most index insurance products rely on mutualization to withstand adverse outcomes. However, the mutualization mechanism requires a sufficiently large number of policyholders. Insufficient demand will weaken the product, beyond simply failing to recover the cost of designing the index.

The purpose of the present paper is to study, through a modeling of the demand for index insurance, the viability of such a product when it competes with a traditional indemnity-based insurance product. The aim here is less about designing an optimal index and more about examining the conditions under which such an index becomes acceptable in a situation where the insurer must meet a solvency requirement. Particular attention will be paid to the impact of the loading factor applied to the premium (this loading factor can be lower in the case of index insurance due to reduced management costs), policyholders' risk aversion, and compensation delays of the competing indemnity-based insurance as levers for achieving this objective. This

analysis will also lead us to propose the construction of a "hybrid" coverage: by hybrid, we mean a combination of traditional indemnity-based insurance with index insurance, which supplements it in certain cases to leverage the best of both approaches, with index insurance not intervening when the basis risk is too high.

The rest of the paper is organized as follows. In Section 2, we set up the notations and formalize the problem of demand in index insurance and its link with solvency. Section 3 then shows conditions to meet (under the assumption that the utility of the policyholder is exponential) in order to achieve a sufficient demand in index insurance. These conditions lead to a natural choice for a hybrid coverage, mixing traditional and index insurance. This hybrid coverage is presented in Section 3.3. A practical example in cyber insurance is then provided in Section 4.

2 Notation and settings

In this section, we formalize the general framework that we consider to study the demand (Section 2.1) and solvency (Section 2.2) of a portfolio of index insurance products—first in the classical case where policyholders are independent, and then by introducing Section 2.3 to add an accumulation component that may arise when a large number of policyholders experience a simultaneous claim.

2.1 Index insurance demand and statement of the problem

We consider a situation where a policyholder has a choice between two insurance products:

- The first one is a "traditional" indemnity-based insurance product, where the loss $Y \geq 0$ of the policyholder in the following year is fully covered.
- The second one is an index insurance product, which is based on $\phi(\mathbf{W})$, where \mathbf{W} is a set of covariates measured after a claim in order to compute the index.

Typically, $\phi(\mathbf{W})$ will be lower than Y , with overcompensation assumed to be rare in index insurance. To simplify, we consider that a policyholder experiences at most one claim; otherwise, \mathbf{W} should be understood as the (different) circumstances of all encountered claims. If there is no claim, $Y = 0$ and $\phi(\mathbf{W}) = 0$.

A traditional way to model the demand in insurance relies on the concept of expected utility, as in Cummins and Mahul [2004], Hao et al. [2018], or Eeckhoudt and Kimball [1992]. A refinement of this approach, especially in the context where there is a choice between different products, is proposed in Braun and Muermann [2004] or Fujii et al. [2016] and relies on the notion of "regret". In this paper, we concentrate on the most classical framework, that is, expected

utility, also because of recent contributions to the design of index-based covers which rely on expected utility maximization (see, for example, Zhang et al. [2019] or Chen et al. [2023]).

Consider a class of utility functions $\mathcal{U} = \{x \in \mathbb{R} \rightarrow U_\alpha(x) : \alpha \in \mathcal{A}\}$, where $\mathcal{A} \subset \mathbb{R}^k$. Every function U_α is assumed to be non-decreasing and strictly concave to represent risk aversion. Each policyholder is associated with a different function, that is, with a different value of α . The decision to buy index insurance compared to indemnity-based insurance is motivated by the maximization of the expected utility of the final wealth.

- In the case of index insurance, the policyholder pays a premium π_ϕ , incurs a loss Y , and receives a compensation $\phi(\mathbf{W})$. The corresponding expected utility is

$$\mathfrak{U}_\phi(\alpha) = E[U_\alpha(\phi(\mathbf{W}) - Y - \pi_\phi)].$$

where $\pi_\phi = (1 + \theta)E[\phi(W)]$, and $\theta > 0$ is the loading factor of the index insurance product.

- In the case of indemnity-based insurance, the difference lies in the fact that the price is π_Y , but the compensation is Y . However, we want to reflect the fact that this compensation is usually paid with a longer delay than with index insurance, which may be problematic for the policyholder who requires liquidity to repair the damages. Therefore, we consider that the compensation will be discounted by a factor $\exp(-\tau)$ for some $\tau \geq 0$. This leads to the following expected utility for this solution:

$$\mathfrak{U}_{Y,\tau}(\alpha) = E[U_\alpha(\{\exp(-\tau) - 1\}Y - \pi_Y)].$$

where $\pi_Y = (1 + \theta_Y)E[Y]$, and $\theta_Y > 0$ is the loading factor of the indemnity-based insurance product.

Let us note that we did not consider the initial capital of the policyholder. This can be taken into account by incorporating it into the parameter α (which may be multivariate). Moreover, let us assume that we consider a function defined on \mathbb{R} to allow the possibility of a negative value for the policyholder's wealth (in which case, a debt is created).

A policyholder with parameter $\alpha \in \mathbb{R}^k$ will choose to rely on index insurance only if

$$\mathfrak{U}_\phi(\alpha) - \mathfrak{U}_{Y,\tau}(\alpha) > 0. \tag{2.1}$$

Here, we assume to simplify that the customer buys one of the two contracts. Alternatively, one could easily consider the option where a third choice of not buying any insurance protection is possible.

2.2 Effect on mutualization

To be viable, an index insurance product needs to be subscribed by a sufficient number of policyholders. Consider a target population of potential customers of large size N , the number of policyholders buying the index insurance contract will be approximately

$$n = N \int \mathbf{1}_{\mathfrak{U}_\phi(\alpha) - \mathfrak{U}_{Y,\tau}(\alpha) > 0} d\mu(\alpha), \quad (2.2)$$

where μ is the cumulative distribution of α among the population.

The question is then to know if this number n is large enough to build a portfolio that is economically viable. Considering that $(Y_i)_{1 \leq i \leq n}$ are the losses of the n policyholders, the loss of the insurance company is

$$L_n(\pi_\phi) = \sum_{i=1}^n \phi(\mathbf{W}_i) - n\pi_\phi.$$

The size of the portfolio should be large enough to ensure that ruin during the next year is avoided with a sufficiently high probability, that is we want

$$\mathbb{P}(L_n(\pi_\phi) \geq 0) \leq \varepsilon, \quad (2.3)$$

where $\varepsilon > 0$ is close to zero.

If we assume that the policyholders are independent and identically distributed, the Central Limit Theorem applies, and

$$n^{1/2} \left\{ \frac{1}{n} \sum_{i=1}^n \phi(\mathbf{W}_i) - \pi_\phi^* \right\} \xrightarrow{n \rightarrow \infty} \mathcal{N}(0, \sigma_\phi^2),$$

where $\sigma_\phi^2 = \text{Var}(\phi(\mathbf{W}))$ and $\pi_\phi^* = E[\phi(\mathbf{W})]$. From this distributional convergence, one can deduce the approximation

$$\mathbb{P}(L_n(\pi_\phi) \geq 0) \approx S \left(\frac{n^{1/2} \theta \pi_\phi^*}{\sigma_\phi} \right),$$

where S is the survival function of a $\mathcal{N}(0, 1)$ variable, and $\pi_\phi = (1 + \theta)\pi_\phi^*$, with $\theta > 0$ being the loading factor of the index insurance product. Hence (2.2) approximately rewrites

$$\frac{n^{1/2} \theta \pi_\phi^*}{\sigma_\phi} \geq S^{-1}(\varepsilon). \quad (2.4)$$

Of course, increasing θ does not necessarily lead to an improvement in solvency, since n decreases with θ as demand is reduced.

We will keep this Gaussian approximation in the next parts, but let us note that this requires the variance of $\phi(\mathbf{W})$ to be finite, which may not be the case for extremely heavy-tailed

distributions. If $\phi(\mathbf{W})$ is extremely heavy-tailed, and if ε is small compared to π_ϕ , other kinds of approximations based on Generalized Pareto distributions may be used, see for example Mikosch and Nagaev [1998] for more details. However, extremely heavy-tailed variables are not our main focus in the present paper, since they are incompatible with the exponential utility approach developed in section 2 (which requires the loss to have a finite Laplace transform).

The independence and identical distribution assumption is justified if the field of application of the index insurance product is restricted to cases where policyholders are not individuals but large geographical areas (such as villages, cities, departments, regions, etc.) or large entities such as companies in the case of cyber insurance, as is the application in this paper. However, despite this restrictive context of application, there remains a slight probability of accumulation, at least on a fraction of the portfolio. For example, in the case of crop insurance (which is probably one of the most famous use cases of index products, see for example Cesarini et al. [2021b] or Barnett and Mahul [2007]), weather events may strike a group of geographical areas simultaneously. To account for that, we propose, in the following section, a simplified way to proceed via the introduction of a shock on the i.i.d. model that materializes the presence of catastrophic events.

2.3 A simplified way to include accumulation phenomena

An accumulation phenomenon occurs when a significant number of policyholders encounter claims in a short period of time. In climate-related risk, this is essentially linked to the proximity between policyholders: policyholders that live in the same area are affected by similar weather conditions. In other cases, the dependence may not solely rely on geographic proximity, such as in cyber insurance. In such cases, it may be hopeless to obtain a clear map of the links between the policyholders. Moreover, trying to model the dependence between policyholders would introduce an additional difficulty in our context, where we take demand into account: even in the case of geographic dependence, this would require modeling a link between the distribution μ (describing the behaviors of the policyholders) and the localization of the potential customers. Calibrating the model would then require an important amount of data that may be difficult to obtain.

Consequently, we consider a simplified case where the accumulation episode materializes via an additional loss, which is heavy-tailed. The total loss of the portfolio is then

$$L_n(\pi_\phi) = A_n + \sum_{i=1}^n \varrho_i \phi(\mathbf{W}_i) - n\pi_\phi,$$

where $\varrho_i = 0$ if policyholder i was part of an accumulation episode, and A_n represents the aggregated amount related to accumulation episodes.

For A_n , we consider a Generalized Pareto distribution, that is

$$\mathbb{P}(A_n \geq t) = \frac{1}{\left(1 + \frac{\gamma t}{ns}\right)^{1/\gamma}},$$

with $\gamma < 1$ and $s > 0$. Here we consider a Generalized Pareto where the scale parameter is proportional to n : this is the idea that the cost of the accumulation episode is proportional to the size of the portfolio.

In this case, the probability of ruin is bounded by

$$\mathbb{P}(L_n(\pi_\phi) \geq 0) \leq \mathbb{P}\left(A_n - \frac{n\theta}{a} \geq 0\right) + \mathbb{P}\left(\sum_{i=1}^n \phi(\mathbf{W}_i) - \left(\frac{1}{\theta} + \left\{1 - \frac{1}{a}\right\}\right) n\theta\pi_\phi^* \geq 0\right),$$

for all $a > 1$. Using the same Gaussian approximation as in section 2.2, the right-hand side is approximately

$$\frac{1}{\left(1 + \frac{\gamma\theta}{as}\right)^{1/\gamma}} + S\left((a-1)\frac{n^{1/2}\theta\pi_\phi^*}{a\sigma_\phi}\right).$$

To make this quantity less than the tolerance ε , we need

$$\frac{n^{1/2}\theta\pi_\phi^*}{\sigma_\phi} \geq \frac{a}{(a-1)} S^{-1}\left(\varepsilon - \frac{1}{\left(1 + \frac{\gamma\theta}{as}\right)^{1/\gamma}}\right), \quad (2.5)$$

if

$$1 < a < \frac{\gamma\theta\varepsilon^\gamma}{s(1-\varepsilon^\gamma)},$$

which imposes that θ should be large enough to absorb the accumulation episode.

Logically, dealing with an additional accumulation risk increases the number of required policyholders to achieve mutualization. Moreover, a constraint appears on the loading factor θ , which should be large enough. Again, since the achievable n tends to decrease when the loading factor increases due to a lower demand, this number may become impossible to reach in some cases.

In the following section, we discuss conditions on the demand for (2.4) and (2.5) to hold in the special case where the utility function is exponential: a simple form of the utility which allows obtaining light constraints on the measure μ .

3 Sufficient conditions for the viability of an index insurance product under exponential utility

In this section, we consider the particular case where the utility function is exponential. This allows us to simplify considerably the formulation of the problem and to provide, in section 3.1, sufficient conditions for an index insurance product to be preferable compared to a traditional one. Then, we consider in section 3.2 the consequences on the solvency of the portfolio. Section 3.3 introduces a way to combine traditional and index insurance to optimize the attractiveness of the product.

3.1 Exponential utility

In this section, we consider $U_\alpha(x) = -(1/\alpha)\exp(-\alpha x)$. The parameter α can be interpreted as a materialization of risk aversion, in the sense that a policyholder with a high value of α will tend to accept a higher premium in exchange for insurance protection against the risk. In this case, the condition (2.1) can be simplified. We introduce the Laplace transform and conditional Laplace transforms of Y ,

$$\begin{aligned}\Psi_Y(\alpha) &= E[\exp(\alpha Y)], \\ \psi_Y(\alpha|\mathbf{w}) &= E[\exp(\alpha Y)|\mathbf{W} = \mathbf{w}].\end{aligned}$$

We assume that $\Psi_Y(\alpha) < \infty$ for all α in the support of μ . Then,

$$\begin{aligned}\mathfrak{U}_\phi(\alpha) - \mathfrak{U}_{Y,\tau}(\alpha) > 0 &\iff -\alpha \{E[\psi_Y(\alpha|\mathbf{W}) \exp(-\alpha\{\phi(\mathbf{W}) - \pi_\phi\})] - \Psi_Y(\alpha') \exp(\alpha\pi_Y)\} > 0, \\ &\iff E[\psi_Y(\alpha|\mathbf{W}) \exp(-\alpha\phi(\mathbf{W}))] < \Psi_Y(\alpha') \exp(\alpha(\pi_Y - \pi_\phi)),\end{aligned}\quad (3.1)$$

where $\alpha' = (1 - \exp(-\tau))\alpha$.

From this expression, we see that condition (2.1) is essentially a matter of bounding the difference $m_Y(\alpha|\mathbf{W}) - \phi(\mathbf{W})$, where $m_Y(\alpha|\mathbf{w}) = \log \psi_Y(\alpha|\mathbf{w})/\alpha$. The difference (which is positive from Jensen's inequality) should be small enough compared to the difference in prices $\pi_Y - \pi_\phi$, and the presence of $\Psi_Y(\alpha')$ allows this difference to go higher when τ increases.

Therefore, a first idea could be to take $\phi(\mathbf{w}) = m_Y(\alpha|\mathbf{w})$ as an indemnity function. But this solution would not be efficient, in the sense that this pay-off would lead to too high a price: from Jensen's inequality, the pure premium would then be

$$E[m_Y(\alpha|\mathbf{w})] > E[Y].$$

Hence, except if we are very close to equality (which would happen only if α is close to zero and/or $Y|\mathbf{W}$ has a variance close to zero), it would become very difficult to offer a premium π_ϕ smaller than π_Y , and in some cases even impossible with a loading factor $\theta_\phi > 0$. In addition to the problem of high prices, this type of contract would typically lead to a too important compensation in many cases. This may collide with some legal constraints depending on insurance regulations¹.

To be compatible with this operational necessity to keep a low price, we consider a pay-off $\phi(\mathbf{w}) = \phi_\beta(\mathbf{w}) = \beta E[Y|\mathbf{W} = \mathbf{w}]$, with $\beta \leq 1$. This choice also allows us to control the probability of over-compensating for a claim. From Chernoff's inequality, this probability is

$$\mathbb{P}(Y - \beta E[Y|\mathbf{W}] < 0 | \mathbf{W} = \mathbf{w}) \leq \psi_Y(\rho|\mathbf{w}) \exp(-\rho(1 - \beta)E[Y|\mathbf{W} = \mathbf{w}]), \quad (3.2)$$

for all $\rho > 0$ such that $\psi_Y(\rho|\mathbf{w}) < \infty$. This inequality can be used to control the proportion of cases where the compensation is too high.

Another approach to define the pay-off is to perform utility maximization as in Zhang et al. [2019] and Chen et al. [2023]. In such an approach, the idea is to maximize a trade-off between large amounts of compensation $\phi(\mathbf{W})$ and an affordable price. However, let us note that, with this approach, $E[\phi(\mathbf{W})]$ may be greater than $E[Y]$ in some situations. Therefore, we prefer in this setting to consider a particular shape of pay-off, determining which difference of price between the indemnity and the index product is acceptable by the customer. If we determine a situation where $\phi(\mathbf{W}) = \beta E[Y|\mathbf{W}]$, associated with a loading factor θ , is an acceptable index product, an additional optimization can be performed if the constraint on $E[\phi(\mathbf{W})] \leq E[Y]$ is not an issue. Finally, we believe that the choice of the value of β is best left in the hands of the insurer, as long as it is clearly stated at the signing of the contract and the policyholder is aware of the structure of the product being purchased. Indeed, one could imagine a product in which an insurer offers various proportions of coverage of the expected average loss, and policyholders choose the proportions corresponding to the premiums they are willing and able to pay. Such an approach, along with its advantages and limitations, is discussed to some extent in Liang et al. [2023], Marone and Sabety [2022], and Woodard [2016].

Proposition 3.1 shows that the index insurance product is chosen by a policyholder with risk aversion α provided that the loading factor θ is small enough and that constraint (3.3) on the conditional Laplace transform at point α holds.

Proposition 3.1 *Assume that*

¹For example, according to French legislation, this could be interpreted as "enrichment without cause", although the existence of a claim may be a protection against this argument. See for example S. Bros, L'assurance paramétrique en assurance de dommages, bjda.fr 2023, Dossier n° 6.

$$\sup_{\mathbf{w} \in \mathcal{W}} \frac{m_Y(\alpha|\mathbf{w}) - \phi_\beta(\mathbf{w})}{E[Y]} < 1 - \beta + \theta_Y. \quad (3.3)$$

Let

$$\eta(\alpha, \beta) = 1 - \beta + \theta_Y - \left\{ \sup_{\mathbf{w} \in \mathcal{W}} \frac{m_Y(\alpha|\mathbf{w}) - \phi_\beta(\mathbf{w})}{E[Y]} \right\}.$$

Then, for $\tau \geq 0$, Condition (2.1) holds if

$$\theta \leq \frac{\eta(\alpha, \beta)}{\beta}.$$

Hence, in this case, there exists an index insurance product with a positive loading factor that is preferable for a policyholder with risk aversion α .

The proof is given in section 6.2.

In this result, condition 3.3 is key and needs to be examined more closely. By Jensen's inequality, $m_Y(\alpha|\mathbf{w}) \geq E[Y|\mathbf{W} = \mathbf{w}]$, so

$$m_Y(\alpha|\mathbf{w}) - \phi_\beta(\mathbf{w}) \geq (1 - \beta)E[Y|\mathbf{W} = \mathbf{w}].$$

Taking the expectation, we see that

$$\frac{E[m_Y(\alpha|\mathbf{w}) - \phi_\beta(\mathbf{w})]}{E[Y]} \geq 1 - \beta. \quad (3.4)$$

Hence, a necessary condition for condition 3.3 is that the difference between the left-hand side and right-hand side of (3.4) is not too high (less than θ_Y). This difference tends to become lower when α tends to zero, confirming the intuition that risk aversion plays against the index insurance product. Moreover, a smaller value of $Var(Y|\mathbf{W} = \mathbf{w})$ will also reduce the gap between $m_Y(\alpha|\mathbf{w}) - \phi_\beta(\mathbf{w})$ and its lower bound $(1 - \beta)E[Y|\mathbf{W} = \mathbf{w}]$. This directly refers to the ability to efficiently predict Y from the available information \mathbf{W} , from which is computed the index.

3.2 Consequences on the solvency of the portfolio

Let us note that proposition 3.1 only provides a sufficient condition for Condition (2.1) to hold. It is valid for all values of τ , including $\tau = 0$, and is then sufficient to obtain a lower bound for the demand, since higher values of τ will increase the disadvantage of the indemnity-based insurance product. In case of $\tau > 0$, there is room for weakening (3.3) and/or the condition on θ .

From proposition 3.1, we easily get the following corollary.

Corollary 3.2 Assume that (3.3) holds for some $\alpha_0 > 0$. Then, if $\theta \leq \eta(\alpha, \beta)\beta^{-1}$, Condition (2.1) holds for all values of $\alpha \in (0, \alpha_0 + h_\beta(\tau))$ with

$$h_\beta(\tau) = F^{-1}(F(\alpha_0) \exp(-\tau) \exp(-\alpha[\pi_Y - \pi_{\phi_\beta}])) - \alpha_0,$$

where $F(\alpha) = \Psi'(\alpha) = E[Y \exp(\alpha Y)]$, and

$$n \geq N\mu((0, \alpha_0 + h_\beta(\tau))). \quad (3.5)$$

This result is a direct consequence of the more general result of Lemma 6.1.

Proposition 3.3 Under the framework of section 2.2 (i.i.d. policyholders and no accumulation phenomenon) and the conditions of Corollary 3.2, condition (2.3) holds for ϕ_β provided that $\theta \leq \eta(\alpha, \beta)\beta^{-1}$ and

$$\eta(\alpha, \beta) \geq \frac{\sigma\beta S^{-1}(\varepsilon)}{N^{1/2}E[Y]\mu((0, \alpha_0 + h_\beta(\tau)))^{1/2}}, \quad (3.6)$$

where

$$\sigma^2 = \text{Var}(E[Y|\mathbf{W}]),$$

which can also be rewritten as

$$\sup_{\mathbf{w} \in \mathcal{W}} \frac{m_Y(\alpha|\mathbf{w}) - \phi_\beta(\mathbf{w})}{E[Y]} \leq 1 - \beta + \theta_Y - \frac{\sigma\beta S^{-1}(\varepsilon)}{N^{1/2}E[Y]\mu((0, \alpha_0 + h_\beta(\tau)))^{1/2}}.$$

In the case where the possibility of accumulation episodes is included as in section 2.3, the result of Proposition 3.4 is slightly modified.

Proposition 3.4 Under the framework of section 2.3 (probability of an accumulation phenomenon described by a Generalized Pareto distribution) and the conditions of Corollary 3.2, condition (2.3) holds for ϕ_β provided that $\theta \leq \eta(\alpha, \beta)\beta^{-1}$ and that, for some $\varepsilon' < \varepsilon$ and $a > 1$,

$$\eta(\alpha, \beta) \geq \max \left(\frac{\sigma\beta a S^{-1}(\varepsilon - \varepsilon')}{(a-1)N^{1/2}E[Y]\mu((0, \alpha_0 + h_\beta(\tau)))^{1/2}}, \frac{a\beta s(1 - \varepsilon'^\gamma)}{\gamma \varepsilon'^\gamma} \right) \quad (3.7)$$

which can also be rewritten as

$$\sup_{\mathbf{w} \in \mathcal{W}} \frac{m_Y(\alpha|\mathbf{w}) - \phi_\beta(\mathbf{w})}{E[Y]} \leq 1 - \beta + \theta_Y - \max \left(\frac{\sigma\beta a S^{-1}(\varepsilon - \varepsilon')}{(a-1)N^{1/2}E[Y]\mu((0, \alpha_0 + h_\beta(\tau)))^{1/2}}, \frac{a\beta s(1 - \varepsilon'^\gamma)}{\gamma \varepsilon'^\gamma} \right).$$

The proofs of these two results are given in section 6.3.

Again, according to Propositions 3.3 and 3.4, solvency can be achieved as long as $(m_Y(\alpha|\mathbf{w}) - \phi_\beta(\mathbf{w}))/E[Y]$ is sufficiently small, which, as we mentioned earlier, can be interpreted as the

ability of \mathbf{W} to capture sufficient information on Y . Let us note that, in this condition, the uniformity with respect to \mathbf{w} is an important weak point: one expects to have situations where it is harder to approximate Y from \mathbf{W} , leading to an increase in $Var(Y|\mathbf{W} = \mathbf{w})$. This leads to the introduction, in the next section, of a hybrid product, mixing indemnity and index insurance, where the use of index insurance is restricted to the most favorable type of events.

3.3 The hybrid product

Let

$$\mathcal{W}_{\alpha}(\epsilon, \beta) = \{\mathbf{w} \in \mathcal{W} : m_Y(\alpha|\mathbf{w}) - \phi_{\beta}(\mathbf{w}) \leq \epsilon\}.$$

We define the following pay-off:

$$\mathfrak{h}_{\alpha, \beta}^{\epsilon}(Y, \mathbf{W}) = \exp(-\tau)Y \mathbf{1}_{\overline{\mathcal{W}_{\alpha, \beta}(\epsilon)}} + \phi_{\beta}(\mathbf{W}) \mathbf{1}_{\mathcal{W}_{\alpha, \beta}(\epsilon)}, \quad (3.8)$$

where $\overline{\mathcal{A}}$ is the complementary set of the set \mathcal{A} .

The idea is that we use index insurance only in cases where we expect this solution to be reliable. From Proposition 3.1, we saw that the unfavorable situations for index insurance compensation are those where $m_Y(\alpha|\mathbf{w}) - \phi_{\beta}(\mathbf{w})$ is large, which motivates the introduction of $\mathcal{W}_{\alpha, \beta}(\epsilon)$. We consider that $m_Y(\alpha|\mathbf{w}) - \phi_{\beta}(\mathbf{w})$ is large if it exceeds a certain threshold ϵ . In this case, policyholders will have a higher utility if they are compensated through indemnity-based insurance (see Section 3.1). As with β , the choice of ϵ is left to the discretion of the insurer. However, the latter should keep in mind that this choice will influence both the proportion of policyholders who prefer compensation through index insurance and the maximum loading factor θ^{\max} at which index insurance can be sold to a given population. We believe that, in practice, the insurer will choose the value of ϵ corresponding to this maximum loading factor, provided that the number of policyholders willing to accept compensation through index insurance is sufficient to make the product economically viable (see Sections 2.2 and 2.3). The impacts of the choice of ϵ in a practical example are discussed in Section 4.3.1.

The premium $\pi_{\mathfrak{h}}$ associated with this product is

$$\pi_{\mathfrak{h}} = (1 + \theta_Y)E \left[Y | \overline{\mathcal{W}_{\alpha, \beta}(\epsilon)} \right] (1 - p_{\epsilon}(\alpha, \beta)) + (1 + \theta)E [\phi_{\beta}(\mathbf{W}) | \mathcal{W}_{\alpha, \beta}(\epsilon)] p_{\epsilon}(\alpha, \beta),$$

where $p_{\epsilon}(\alpha, \beta) = \mathbb{P}(\mathbf{W} \in \mathcal{W}_{\alpha, \beta}(\epsilon))$. We apply the same loading factor θ_Y as for indemnity insurance to cases where exact compensation is offered, the lower loading factor θ being applied only to the index part. The higher loading factor of indemnity insurance directly reflects the increased claims management costs associated with this type of coverage. This hybrid design bears a conceptual resemblance to the blended insurance scheme proposed by Zhang [2024].

Proposition 3.5 provides a condition for the hybrid product $\mathfrak{h}_{\alpha, \beta}^{\epsilon}$ to be chosen instead of the traditional indemnity-based contract.

Proposition 3.5 *If*

$$\eta_{\epsilon}(\alpha, \beta) = 1 - \beta + \theta_Y - \frac{\epsilon}{E[Y|\mathcal{W}_{\epsilon}(\alpha, \beta)]p_{\epsilon}(\alpha, \beta)}, \quad (3.9)$$

then, if $\theta \leq \eta_{\epsilon}\beta^{-1}$, the policyholder with risk aversion less than α prefers the contract defined by the pay-off $\mathfrak{h}_{\alpha, \beta}^{\epsilon}$ for all $\tau \geq 0$.

From the fact that (3.9) should be non-negative, we see that the set $\mathcal{W}_{\alpha}(\epsilon, \beta)$ should not be too small, otherwise the probability $p_{\epsilon}(\alpha, \beta)$ could make the left-hand side larger than $1 - \beta + \theta_Y$. On the other hand, one could be tempted to take a low value for ϵ to control the difference between $m_Y(\alpha|\mathbf{w})$ and $\phi_{\beta}(\mathbf{w})$, but this mechanically tends to make $\mathcal{W}_{\epsilon}(\alpha, \beta)$ shrink. Let us also note that a too important increase of ϵ introduces more constraints on θ : a decrease of $\eta_{\epsilon}(\alpha, \beta)$ makes condition (3.9) easier to achieve, but the loading factor θ then should be smaller.

Proof. Similarly to the case of a full index product, the situation $\mathfrak{U}_{\mathfrak{h}_{\alpha, \beta}^{\epsilon}}(\alpha) - \mathfrak{U}_{Y, \tau}(\alpha) > 0$ implies

$$E[\alpha \exp(Y - \mathfrak{h}_{\alpha, \beta}^{\epsilon}(Y, \mathbf{W}))] \leq E[\exp(\alpha'Y)] \exp(\alpha[\pi_Y - \pi_{\mathfrak{h}}]). \quad (3.10)$$

With $\alpha' = (1 - \exp(-\tau))\alpha$

The left-hand side rewrites

$$\begin{aligned} & (1 - p_{\epsilon}(\alpha, \beta))E[\exp(\alpha'Y)] + E\left[\exp(\alpha\{m_Y(\alpha|\mathbf{W}) - \phi_{\beta}(\mathbf{W})\}) \mathbf{1}_{\mathbf{W} \in \mathcal{W}_{\alpha, \beta}(\epsilon)}\right] \\ & \leq (1 - p_{\epsilon}(\alpha, \beta))E[\exp(\alpha'Y)] + \exp(\alpha\epsilon)p_{\epsilon}(\alpha, \beta) \\ & \leq E[\exp(\alpha'Y)] \left\{ (1 - p_{\epsilon}(\alpha, \beta)) + \frac{\exp(\alpha\epsilon)}{E[\exp(\alpha'Y)]} p_{\epsilon}(\alpha, \beta) \right\} \\ & \leq E[\exp(\alpha'Y)] \exp(\alpha\epsilon). \end{aligned}$$

Moreover,

$$\pi_Y - \pi_{\mathfrak{h}} = (1 + \theta_Y - \beta - \beta\theta)E[Y|\mathcal{W}_{\alpha, \beta}(\epsilon)]p_{\epsilon}(\alpha, \beta).$$

Hence, a sufficient condition for (3.10) is

$$\frac{\epsilon}{E[Y|\mathcal{W}_{\alpha, \beta}(\epsilon)]p_{\epsilon}(\alpha, \beta)} \leq 1 - \beta + \theta_Y - \beta\theta,$$

which means that θ should be less than $\eta_{\epsilon}(\alpha, \beta)\beta^{-1}$. As for the proof of Proposition 3.1, $\mathfrak{U}_{\mathfrak{h}_{\alpha, \beta}^{\epsilon}}(\alpha) - \mathfrak{U}_{Y, \tau}(\alpha) > 0$ implies that $\mathfrak{U}_{\mathfrak{h}_{\alpha, \beta}^{\epsilon}}(\tilde{\alpha}) - \mathfrak{U}_{Y, \tau}(\tilde{\alpha}) > 0$ for any $\tilde{\alpha} \leq \alpha$. ■

4 Practical illustration

We illustrate in this section the results obtained previously with an example in the field of cyber insurance. We begin by studying the solvency requirements discussed in sections 2.2, 2.3 and

3.2, and we end by showing how to deduce in which cases index insurance is relevant (see section 3.3). In section 4.1, we describe the database and the choices made for the different parameters required in our models. Section 4.2 focuses on the requirements of index insurance to ensure solvency, and finally, section 4.3 is devoted to the design of the hybrid product presented in section 3.3.

4.1 Description of the context and the database

One of the complexities of cyber insurance lies in the difficulty of assessing the economic losses associated with a cyberattack. These losses are multifaceted and cannot be reduced to just damage to digital assets. Numerous high-profile examples, such as the Wannacry attack or the Colonial Pipeline breach, show that the damages are less related to the ransom amount than to the business interruption.

The database.

We consider a synthetic database of business interruptions related to cyber attacks². The events contained in this database reflect real-life situations, but random noise has been introduced for the sake of confidentiality, and the data have been rescaled to reflect market trends which are consistent with the premium levels we consider in the following.

The database contains $n = 10,000$ claims. On each claim i , the following information is available:

- the loss amount Y_i ;
- the duration of the business interruption T_i ;
- $X_i \in \{t_1, \dots, t_5\}$ is a qualitative variable related to the type of service that is impacted ;
- an indicator variable $\delta_i \in \{0, 1\}$, equal to 1 if the policyholder was able to activate a backup plan that potentially mitigated the impact of the business interruption;
- an indication of the quality of the backup plan $B_i \in (0, 1)$, corresponding to the expected proportion of activity that is preserved once the backup plan has been triggered ; this quantity is anticipated from security audit performed before the claim ;
- $\Lambda_i = (T_i - U_i)_+$ where U_i is the delay before triggering the backup plan.

The available information just after the claim is $\mathbf{W}_i = (T_i, X_i, \delta_i, B_i, \Lambda_i)$.

Some descriptive statistics are provided in Table 1. We see that in the absence of a backup plan, policyholders are on average exposed to a loss nearly 6,000 euros higher than those with a backup plan. This difference is even more striking when considering that the higher losses in the

²Available here: https://github.com/dnkameni/cyber_cloud_interruption/blob/main/loss_data.csv

absence of a backup occur over significantly shorter average interruption durations. For further analysis, see Figure 8 in the appendix, which presents the densities of these two variables.

Variable	Mean	Minimum	Maximum	Standard deviation
Y (in 10^3 euros)	109.08	1.09	2128.59	113.37
T (in days)	2.09	0.01	19.65	2.39
$Y, \delta = 1$	95.56	1.09	902.21	85.31
$T, \delta = 1$	3.20	0.01	16.09	2.64
$Y, \delta = 0$	118.25	2.01	2128.59	128.16
$T, \delta = 0$	1.33	0.01	19.65	1.85

Table 1: Descriptive statistics for the database.

The duration of business interruption T is a key factor in assessing the consequences of cyber attacks, as highlighted in several studies, including Tam et al. [2023]. Indeed, this variable is expected to be strongly correlated with the amount of loss. In our dataset, the linear correlation coefficient between losses and duration of business interruption is 0.57. This value is 0.65 for policyholders who activated a backup plan and 0.75 for policyholders who failed to activate a backup plan. Figure 1, which shows a plot of losses due to business interruptions against the duration of the interruptions, appears to confirm these observations. Moreover, the duration of a business interruption can be measured immediately after the occurrence of the interruption, without requiring lengthy and costly expert evaluations. It is therefore quite natural to use this variable as the core component for constructing an index-based insurance product. Generally, the impact of a given duration of service interruption depends heavily on the sector of activity and the size of the affected company. In this simplified example, we consider a portfolio composed of policyholders with homogeneous profiles — that is, operating in the same sector, located in the same country, and of similar size.

Premium amount.

In this simplified framework, we assume that each policyholder pays the same premium, calculated as the average compensation within the portfolio. To determine a reasonable value for the loading factor θ_Y , we refer to the loss ratios observed on the market between 2020 and 2022 for medium-sized companies (according to the terminology used in the LUCY — Light Upon Cyber Insurance — report on the French market, published by AMRAE³). The reported loss ratios are 45% (2020), 36% (2021), and 100% (2022), yielding an average of approximately 60%. For the purpose of this illustration we thus make the assumption that the premium π_Y is calibrated to yield an average margin of 40% above the pure premium, that is, we set $\theta_Y = 0.4$.

³Association pour le Management des Risques et l'Assurance de l'Entreprise. See the AMRAE report LUCY, <https://www.amrae.fr/bibliotheque-de-amrae/lucy-light-upon-cyber-insurance-2024-edition>




Figure 1: Losses Y against duration of interruption T with a distinction between policyholders whose backup plans were activated and those whose backup plans failed to kick-in. We observe an increasing relationship between losses and duration of interruption

It is important to emphasize that this loading factor is not intended to accurately reflect actual market practices: premium levels can vary significantly from one customer to another, and loss ratios are not always anticipated. The current state of the cyber insurance market is still highly unstable. As shown in the same LUCY report, the amount of premiums has fluctuated considerably from year to year, illustrating the constant re-evaluation of risk by insurers in response to their prior underwriting results.

Frequency of claims.

The dataset we consider provides information about the severity of claims, but not about their frequency, as there is no indication of exposure. Thus, it informs us about the pay-off in the event of a claim, but not about the probability of such a claim occurring.

To assume a plausible value for frequency, we set the probability of experiencing an incident to $p = 0.06$. This value is inspired by the LUCY report (492 medium-sized companies in the sample, with 30 claims reported). Once again, we recall that, for simplification, our framework assumes that a policyholder does not experience more than one claim per year.

Demand and risk aversion.

In this application, we adopt an exponential utility function and must specify a measure μ to describe the distribution of risk aversion among potential policyholders. In the absence of a rigorous market study to assess price elasticity, modeling the distribution μ is the most questionable assumption. Again, the goal here is only to illustrate the methodology presented in this paper, not to provide a reliable estimate of cyber insurance demand.

We consider here that μ is a shifted exponential distribution, that is

$$d\mu(t) = \lambda \exp(-\lambda(t - \alpha_-)) \mathbf{1}_{t \geq \alpha_-}.$$

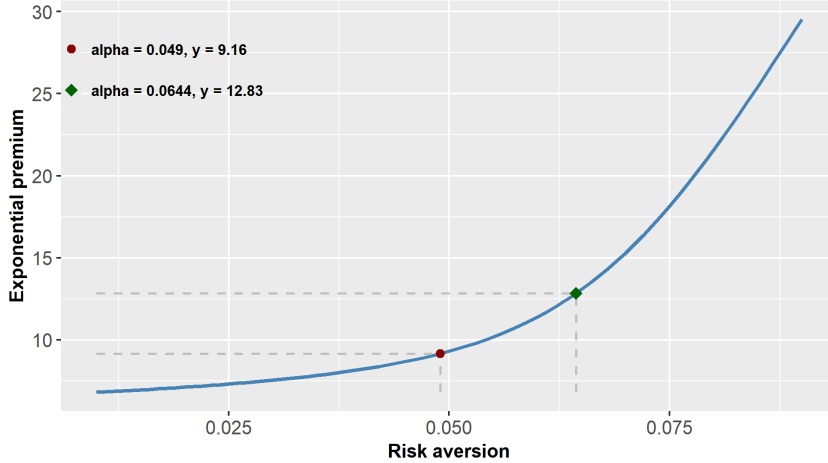


Figure 2: Estimation of $\alpha \rightarrow \log \Psi_Y(\alpha)/\alpha$ (exponential premium). The left point corresponds to the value of risk aversion corresponding to an exponential premium equal to π_Y . The right point (blue) corresponds to the case where π_Y increases by 40%.

To determine the value of α_- , we observe that the target population of policyholders who already subscribed an insurance contract accepted a price π_Y . Since their preferences are described via an exponential utility, this is possible only if their risk aversion α is high enough. If we do not take into account the potential discount factor τ at this stage, this means that

$$\frac{\log \Psi_Y(\alpha_-)}{\alpha_-} = \pi_Y.$$

Estimating empirically $\alpha \rightarrow \Psi_Y(\alpha)$ from the database (see Figure 2), we get $\alpha_- = 0.049$.

Next, to consider a proper value for λ , we make the assumption that half of the population of the policyholders is ready to accept an increase of 40% of the premium. This choice is arbitrary, but is motivated by the fact the LUCY report noted an increase of 84% of the collected premiums over 2022, for an increase of 53% of the number of policyholders in the perimeter of the study (while deductible increase and insurance capacity stays stable). This seems to indicate that all of the current policyholders were ready to accept an increase of approximately 20% of their premium. The proportion 40% that is taken to set the value of λ is then based on twice this number. This leads to $\lambda = 45.08$.

4.2 Empirical analysis of the solvency requirements of index insurance

In this section, we assume that a given population is exposed to an insurance market offering two types of insurance contracts: index insurance and indemnity-based insurance. We also assume that individuals are rational and will choose the insurance product that they expect to have the highest positive impact on their financial well-being. In other words, they will select the insurance product that maximizes their expected utility of wealth (see Section 2.1). We

then study how the number of individuals in this target population who accept index insurance varies with certain factors that are likely to influence the demand for this type of insurance. This number is compared to the minimum number of policyholders needed to ensure, with high probability, the solvency of an index insurer in both classical situations and situations involving accumulation of losses (see Sections 2.2 and 2.3).

We set the tolerance level ε at 0.5%. This choice is motivated by the regulatory requirement in Europe that insurers must hold sufficient capital to ensure a 99.5% probability of remaining solvent over a one-year period (see Scherer and Stahl [2021]). The values of the parameters used in the accumulation setting (see Section 2.3) are set at $s = 0.003$, $\gamma = 0.5$, and $a = 2.4$. These values are chosen such that the condition $1 < a < \frac{\gamma\theta\varepsilon^\gamma}{s(1-\varepsilon^\gamma)}$ holds for $\theta > 0.18$ (recall that $\theta_Y = 0.4$).

To construct $\phi_\beta(\mathbf{W}) = \beta E[Y|\mathbf{W}]$, we set $\beta = 0.9$ and estimate $E[Y|\mathbf{W}]$ using four different models: a linear model (an approach similar to that of Giné et al. [2007]), a regression tree model, a random forest model, and an eXtreme Gradient Boosting (XGBoost) model. The performances of these models on the in-sample data are presented in Table 2. The random forest and XGBoost models exhibit better performance than the linear and regression tree models. This indicates that these models have a higher probability of accurately predicting the losses suffered by policyholders for a given index value.

Model	<i>RMSE</i>	R^2	<i>MAE</i>	<i>Correlation</i>
Linear model	4.307	0.599	2.813	0.774
Regression trees	4.197	0.619	2.787	0.787
Radom forests	3.018	0.811	1.897	0.901
XGBoost	2.601	0.856	1.741	0.925

Table 2: Models used to build the index payout ϕ .

These high accuracies support the use of index insurance, as shown in Figure 3. This figure presents a plot of the number of policyholders who accept index insurance in the target population against the delay in compensation of the competing indemnity-based insurance product. We observe that, for a given compensation delay in the indemnity-based product, index insurance products whose payouts are built using random forest or XGBoost models are likely to attract more policyholders. This is explained by the fact that the superior performance of these models reduces basis risk, which is one of the major challenges of index insurance. The reduced basis risk increases the satisfaction policyholders derive from index insurance, thereby enhancing the product’s attractiveness. This highlights the positive impact that machine learning can have on the design and commercialization of index insurance. However, the less performant linear and regression tree models have the advantage of being more interpretable and easier to explain to

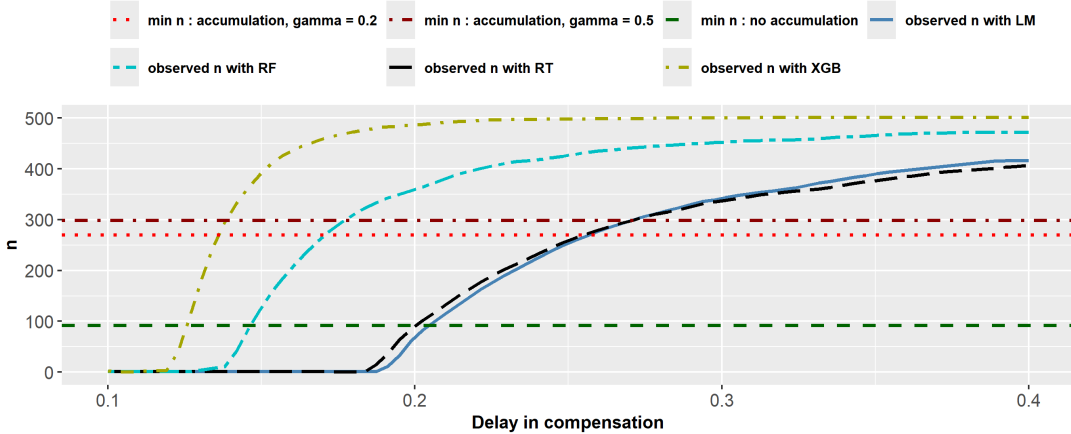


Figure 3: Number of policyholders n in the population of size N who prefer index insurance, as a function of the delay in compensation τ of the competing indemnity-based insurance product. The values of n are represented for four index payout models. The reference values are computed for $\varepsilon = 0.5\%$

policyholders and regulators.

The results of the linear regression presented in Table 3 are consistent with what is typically expected in practice. Specifically, we observe that the losses suffered by policyholders increase with longer interruption times. Conversely, these losses tend to decrease in the presence of a backup plan, with longer durations of backup use, and with higher backup efficiency.

Constructing $\phi_\beta(\mathbf{W})$ is equivalent to learning the relationship between the loss Y and the index \mathbf{W} . For this, we use the entire dataset described in Section 4.1. However, we consider a target population of size $N = 500$, which approximately corresponds to the number of policyholders in the LUCY study (medium-sized companies). Figure 4 shows how the number of policyholders n in the population N who accept index insurance varies with the mean risk aversion in the population $\bar{\alpha}$, and the loading factor of index insurance θ . This figure also illustrates the minimum number of policyholders required in the portfolio to ensure solvency with a 99.5% probability under both classical and accumulation scenarios.

As expected, the minimum number of policyholders in an index insurance portfolio required to ensure solvency with a probability of 99.5% in the presence of accumulation is higher than the number needed when there is no accumulation. This is because greater mutualization is necessary to address accumulation and to efficiently "dilute" large losses among policyholders. For the same reason, in the presence of accumulation, this minimum required number of policyholders increases with the heaviness of the tails of the distribution of accumulation events (see Section 2.3), as shown in Figures 3 and 4.

Panels (a) and (b) of Figure 4 empirically confirm the intuition that risk aversion works

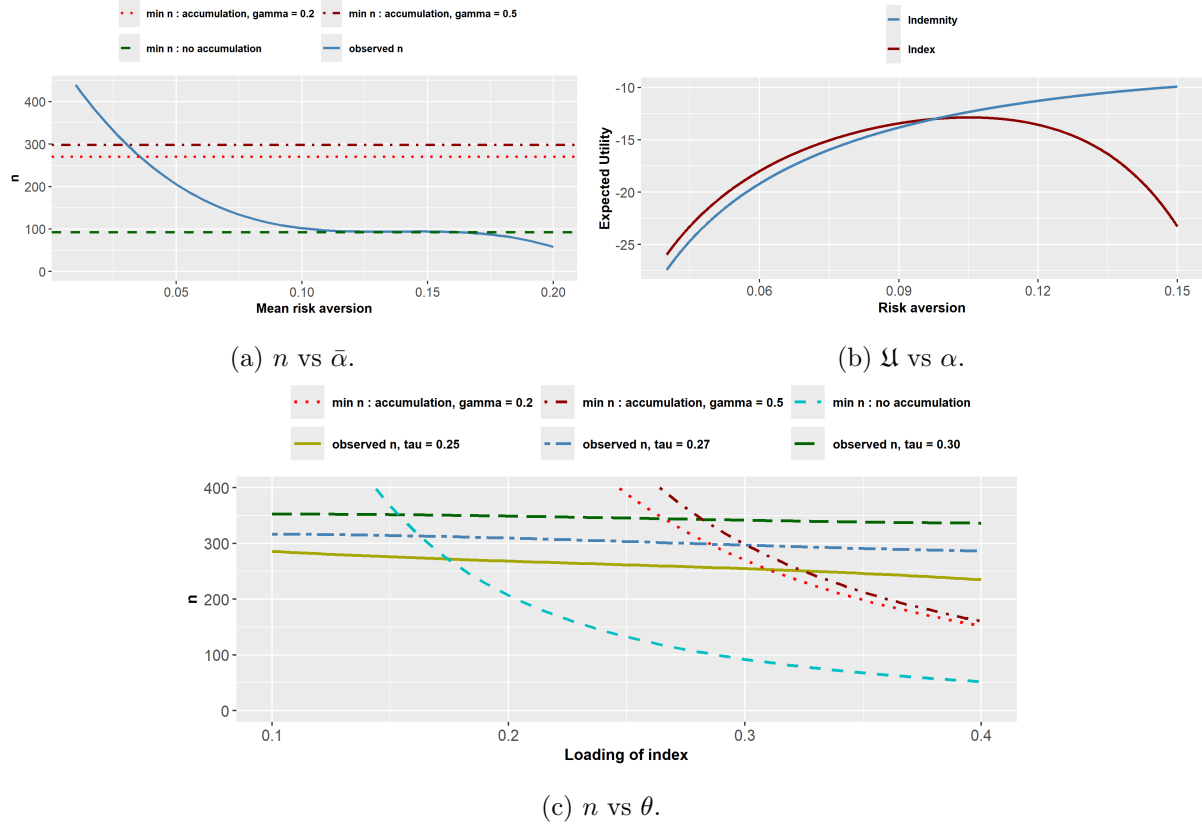


Figure 4: Number of policyholders n in the population of size N who prefer index insurance, as a function of the mean risk aversion in the population $\bar{\alpha}$ and the loading factor of index insurance θ . The reference values are computed for $\varepsilon = 0.5\%$. Panel (b) shows a plot of expected utility against risk aversion for index and indemnity-based insurance.

against index insurance. This conclusion aligns with that of Clarke [2016], which states that the optimal demand for index insurance at any positive price is zero for infinitely risk-averse individuals. Panel (a) shows a steady decrease in the number of policyholders who accept index insurance as the mean risk aversion in the population increases, all else being equal. This results from the fact that, above a certain level of risk aversion, the utility policyholders derive from index insurance drops below that of indemnity insurance and continues to decrease rapidly with increasing risk aversion, as shown in panel (b). According to Clarke [2016], this can be explained by the presence of basis risk⁴ in index insurance. Indeed, basis risk can reduce the wealth of policyholders (on average), thereby lowering their expected well-being or expected utility. A risk-averse individual may therefore choose indemnity-based insurance instead, which in this setting provides higher compensation. However, moderately risk-averse individuals might behave

⁴This is the risk that the payout of the index insurance contract differs from the actual loss suffered by policyholders Clement et al. [2018]

differently, considering that the lower price of index insurance is sufficient to justify the basis risk. This distinct behavior of moderately risk-averse individuals is also highlighted by Clarke [2016]. Finally, note that this analysis holds if the delay in compensation of indemnity-based insurance is moderate. If the competing indemnity product on the market has a significantly longer compensation delay, then even extremely risk-averse individuals might opt for index insurance, as the delay in indemnity-based compensation would considerably reduce the expected present value of their wealth.

This observation appears to be confirmed by Figure 3, which shows the relationship between the number of policyholders who prefer index insurance, n , and the delay in compensation of the competing indemnity-based insurance product, τ , for a fixed distribution of risk aversion and a fixed index insurance loading factor. We observe that it might be impossible to launch an index insurance product on the market if it does not efficiently address the issue of delays in indemnity-based compensation. Specifically, if the difference in compensation speed between the two products is not sufficiently high in favor of index insurance, then the number of individuals in the population who accept index insurance will be insufficient to ensure solvency with a 99.5% probability, regardless of whether accumulation is present. In practice, a low difference in compensation time could result from excessive paperwork or verification procedures imposed by a regulator on index insurance, to the extent that it loses its speed advantage.

The number of individuals in the population who prefer index insurance decreases with an increase in the loading factor θ of the index insurance contract, all else being equal. This is illustrated in Panel (c) of Figure 4, and it is a result that was expected. Indeed, as the loading factor of index insurance increases, policyholders are presented with a product that not only suffers from basis risk but is also priced closer to an equivalent indemnity-based insurance contract ($\theta_Y = 0.4$). Some individuals will then prefer to pay the same price for the indemnity-based product, which offers better coverage, assuming the compensation speed is reasonable. In practice, the lower loading factor of index insurance compared to its indemnity-based counterpart is achieved by eliminating claims management and expert assessment costs.

We also observe in panel (c) a decrease in the minimum number of policyholders required to ensure solvency at a 99.5% probability as the index insurance loading factor increases. This is because, as index insurance becomes more expensive, the insurer earns more income from each policy sold. This increased income allows the insurer to cover all claims with a smaller portfolio, whether under classical loss or accumulation scenarios. This decreasing relationship between the solvency thresholds and θ enables the insurer to determine a minimum loading factor $\theta^{\min} \approx 0.18$ at which the index insurance product can be sold competitively relative to indemnity-based insurance (with $\theta_Y = 0.4$). If an insurer were to offer the same index insurance product designed in this section to a population with the same characteristics as ours, applying a loading factor of 0.18 would result in a portfolio of slightly more than 260 policyholders, which

would be sufficient to ensure solvency with a 99.5% probability in the absence of accumulation. Attempting to increase the number of policyholders by reducing the premium loading could expose the insurer to the risk of insolvency. An alternative solution could be to optimize claims management procedures (e.g., using Artificial Intelligence) to accelerate them, thereby increasing demand for the index insurance product as shown for different values of τ in panel (c). Another solution could be to educate the population about the advantages of index insurance, aiming to mitigate the influence of risk aversion on demand (see Panel (a)).

4.3 Empirical analysis of a hybrid approach

In this section, we consider an insurance setting where an insurer can compensate losses using either an index insurance mechanism or an indemnity-based insurance mechanism. This insurer aims to propose each type of coverage where it is most suitable for policyholders and therefore more likely to be accepted. The identification of the various cases uses a criterion which is built using information on both the behavior and preferences of policyholders, as well as data on the index. However, the ultimate goal is to rely solely on index values to determine the type of coverage to apply. This section is an application of Section 3.3.

Identifying the cases where index insurance is more suitable for compensation amounts to determining the values of the index \mathbf{w} who belong to the set

$$\mathcal{W}_\alpha(\epsilon, \beta) = \{\mathbf{w} \in \mathcal{W} : m_Y(\alpha|\mathbf{w}) - \phi_\beta(\mathbf{w}) \leq \epsilon\},$$

for given values of ϵ and β . This set corresponds to cases where the utility of wealth policyholders derive from index insurance compensation exceeds that which they derive from indemnity-based insurance compensation. To estimate the quantity $\Delta(\mathbf{w}) = m_Y(\alpha|\mathbf{w}) - \phi_\beta(\mathbf{w})$ —and thus determine whether \mathbf{w} belongs to the set $\mathcal{W}_\alpha(\epsilon, \beta)$ —we use two approaches:

- regression tree models (see, for example, Breiman [2017] or Loh [2014]) for their simplicity and interpretability;
- eXtreme Gradient Boosting (XGBoost) models for their good performances in predicting losses on our dataset (see section 4.2).

The choice of the model to use in practice will depend on the trade-off one is willing to make between model interpretability and model performance.

4.3.1 Discussion on ϵ and methodology of analysis.

Recall from Proposition 3.5 that setting a value for ϵ indirectly determines a maximum value $\theta^{\max} = \eta_\epsilon(\alpha, \beta)\beta^{-1}$ for the index insurance loading factor θ . Additionally, the value of ϵ influences the number or proportion ($p_\epsilon(\alpha, \beta)$) of compensations for which index insurance is

preferable, as it directly affects the size of the set $\mathcal{W}_\alpha(\epsilon, \beta)$. This dual influence of ϵ is illustrated in Figure 5, using the same dataset described in Section 4.1.

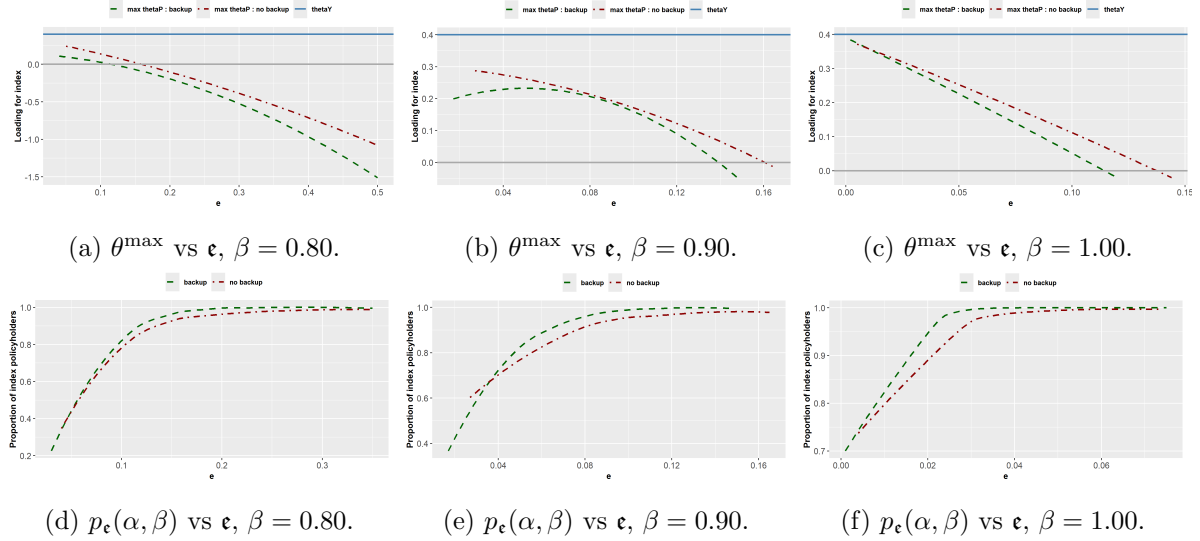


Figure 5: Maximum loading factor of index insurance $\theta^{\max} = \eta_\epsilon(\alpha, \beta)\beta^{-1}$ and proportion of compensations for which index insurance is preferable $p_\epsilon(\alpha, \beta)$ as functions of the parameter ϵ for various values of β . The value of θ_Y is also plotted in panels (a), (b), and (c) for comparison.

Figure 5 shows that increasing the value of ϵ reduces the maximum loading factor θ^{\max} of index insurance and thus lowers the price at which the index insurance contract can be sold to policyholders. Conversely, increasing ϵ raises the proportion of compensations for which index insurance is more suitable. A direct explanation for this is that a higher value of ϵ makes the condition $m_Y(\alpha|\mathbf{w}) - \phi_\beta(\mathbf{w}) \leq \epsilon$ easier to satisfy, thereby increasing the size of the set $\mathcal{W}_\alpha(\epsilon, \beta)$.

In practice, this means that when choosing ϵ , an insurer may need to balance between the proportion of index insurance compensations in their portfolio and the maximum loading that can be applied to the index insurance product. Indeed, Figure 5 empirically demonstrates that increasing the share of index insurance compensation in a portfolio implies the necessity of offering the index insurance component of the hybrid contract at a lower price.

The final choice of ϵ will be made following a solvency and ruin analysis, similar to that conducted in Section 4.2. This analysis will help the insurer select the value of ϵ that maximizes revenue while minimizing the probability of ruin. The existence of such an optimal value is expected due to the dual and opposing effects of ϵ on both the maximum loading of index insurance and the proportion of compensations for which index insurance is preferable. Other advantages of index insurance, such as faster compensation (which can foster client loyalty), reduced volatility, and simplified claims management, could also be factored into the insurer's decision regarding the final proportion of index insurance compensations in their portfolio.

With this in mind, we propose Algorithm 1 to identify the cases where compensation via index insurance is the best option.

Algorithm 1 Identification of compensation preferences

- 1: $\alpha \leftarrow$ estimate of a target risk aversion¹
- 2: $\beta \leftarrow$ a suitable value to control overcompensation
- 3: Set an optimal ϵ^* based on $\theta^{\max}(\alpha, \beta, \epsilon)$ and $p_{\epsilon}(\alpha, \beta)$ ²
- 4: Compute $\theta^{\max}(\alpha, \beta, \epsilon^*)$ and $p_{\epsilon^*}(\alpha, \beta)$
- 5: Train a model (regression trees and XGBoost in our case) to predict:

$$\Delta(\mathbf{w}) = m_Y(\alpha|\mathbf{w}) - \phi_{\beta}(\mathbf{w})$$

- 6: Flag all cases (leaves for regression trees and individuals for XGBoost) where $\sup_{\mathbf{w} \in \mathcal{W}} \Delta(\mathbf{w}) \leq \epsilon^*$ as suitable for index insurance and the rest as suitable for indemnity-based insurance
-

Applying this algorithm gives the insurer a tool to identify the cases where compensation via index insurance is more suitable. These cases constitute a proportion $p_{\epsilon^*}(\alpha, \beta)$ of all compensations in the portfolio. Furthermore, the insurer has an indication on the maximum loading factor that can be applied to the index insurance component of the proposed hybrid product, given by $\theta^{\max}(\alpha, \beta, \epsilon^*)$.

Regarding the choice of β , Figure 5 shows that low values of β , although advantageous in considerably reducing the risk of overcompensation, also decrease the attractiveness of index insurance compensation. Indeed, panels (a) and (d) of Figure 5 show that for $\beta = 0.8$, if the insurer applies the maximum possible loading to index insurance, only 20% of the portfolio will prefer index insurance in the presence of a backup plan. This proportion rises to 70% when $\beta = 1$ (see panels (c) and (f) of Figure 5). This can be explained by the fact that for very low values of β , the reduction in the price of the index insurance product fails to offset the decrease in utility or satisfaction experienced by policyholders due to the corresponding reduction in index insurance compensation. As a result, indemnity-based insurance compensation appears to be a more attractive option. In the rest of this section, we set $\beta = 0.9$. Since the efficiency of the backup plan plays a particular role in the severity of losses, we distinguish between claims associated with no backup plan (or a backup plan that was not triggered quickly enough, $\delta = 0$) and those where the backup program successfully reduced the impact of the incident ($\delta = 1$).

¹According to proposition 3.5, if $\theta \leq \theta^{\max}$, then all policyholders with risk aversion less than α will prefer the proposed hybrid contract compared to a full indemnity-based contract.

²The optimal value of ϵ could be set by considering solvency requirements (see sections 2.2 and 2.3), profitability targets, and expert recommendations on the desired structure of the insurance portfolio (see figure 5).

4.3.2 Applications with regression trees and XGBoost models.

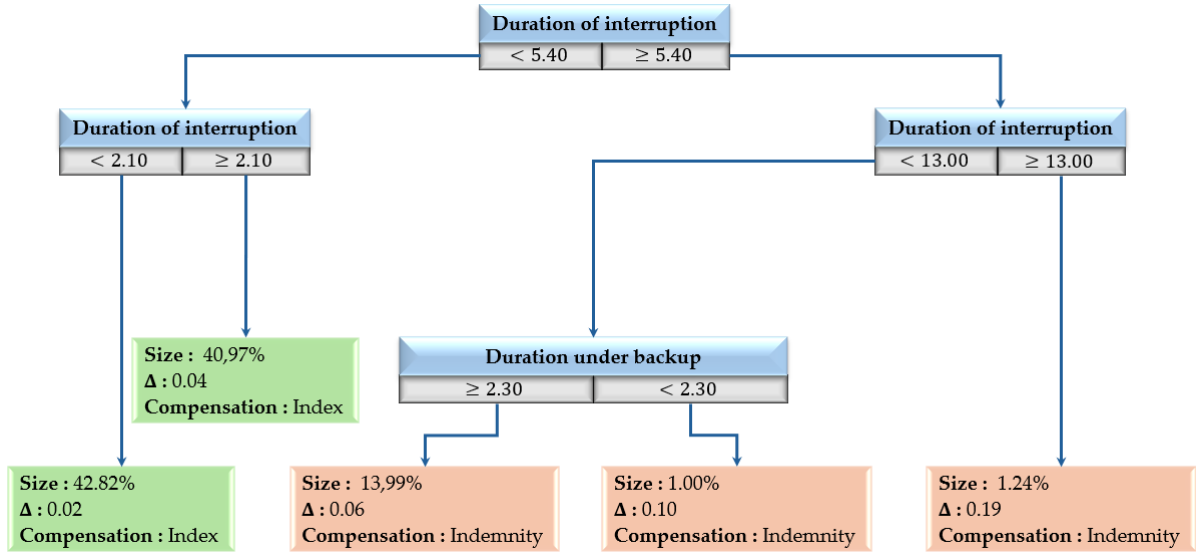
In these applications, the values of ϵ^* are chosen arbitrarily to illustrate our methodology and the proposed hybrid product. Figure 6 and figure 7 display the two regression trees and the two scatter plots obtained in our setting for the regression trees and the XGBoost models respectively.

Applying algorithm 1 with regression trees yields a maximum loading of the index insurance part of the contract (θ^{\max}) equal to 0.32 in cases where the backup plan is triggered on time, and 0.26 in cases where the backup plan fails or is absent. These values are respectively 0.27 and 0.22 when algorithm 1 is used with an XGBoost model. Recall that the loading factor for the indemnity-based part of the contract is 0.40. When regression trees are used, our methodology reveals that compensation with index insurance is preferable in 83.79% of cases, when the backup plan is triggered fast enough. This proportion is 67.87% in situations where the backup plan fails or is absent. When an XGBoost model is used, the previous proportions are respectively 89.43% and 49.60%. These percentages correspond to the shares of policyholders who suffered a loss in the portfolio who are likely to prefer compensation via index insurance.

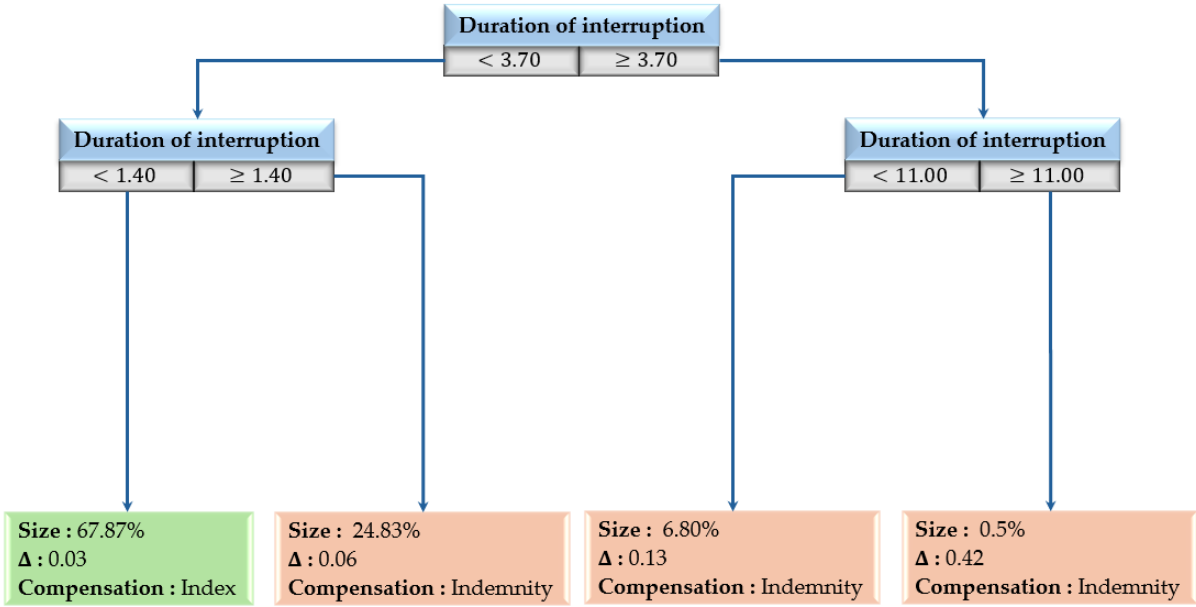
Note that the XGBoost model allows for individual identification of the preferred or ideal compensation types, unlike the regression tree models, which perform a clustering of compensations and determine the ideal compensation type per cluster. This difference in approach, along with the difference in predictive performance between the two models, are among the reasons that could justify the differences in the values of θ^{\max} and $p_{\epsilon^*}(\alpha, \beta)$ obtained from the two models.

The results presented in figure 6 and in figure 7 support the use of index insurance compensation for short durations of business interruption and, consequently, for small losses (as clearly shown in panel (a) of figure 7). This is a favorable outcome, as index insurance tends to exhibit lower basis risk for small losses. Indeed, Lopez and Thomas [2023] show that basis risk generally increases with the size of the loss. In other words, our methodology recommends applying index insurance in contexts where it performs best. Combined with the other advantages of index insurance, this may explain why policyholders in such situations show a preference for it.

Alternative segmentation techniques could also be used for this identification task. Artificial Intelligence could be incorporated to either improve the efficiency of the choice of ϵ^* or enhance the quality of the segmentation of compensations. Once the cases in which a preference for compensation with index insurance are identified, a more detailed analysis of the claims and policyholders' profiles within each preference group could be conducted to build an index insurance payout model tailored to each group where it is preferred. This approach could help reduce overall basis risk and improve the quality of the proposed hybrid insurance contract. Authors such as Hernández-Rojas et al. [2023], Lin et al. [2023], and Cesarini et al. [2021a] also highlight the significant contributions that Artificial Intelligence and Machine Learning could bring to the



(a) Backup plan was successful ($\delta = 1$).



(b) Backup plan failed ($\delta = 0$).

Figure 6: Regression trees used to cluster compensations according to preferences in compensation type (index or indemnity). For each leaf, information is provided on the proportion of compensations in that leaf (size), the value of $\Delta(\mathbf{w})$ for that leaf, and the decision (for the value of \mathbf{e}^* chosen).

development of index insurance payout models.

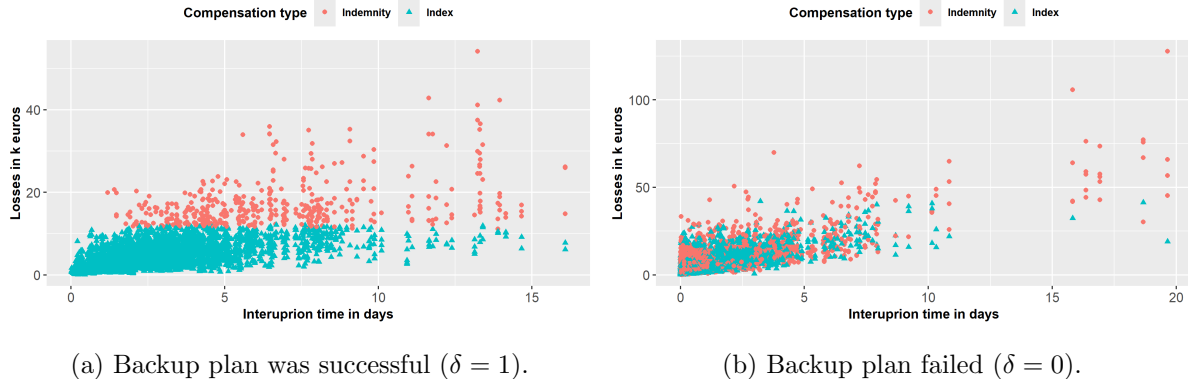


Figure 7: Ideal compensation type obtained by applying algorithm 1 with XGBoost models. In this case, $\Delta(\mathbf{w})$ is calculated at an individual level and not in clusters as is the case with regression trees models.

5 Conclusion

In this paper, we proposed a framework to analyze the introduction of an index insurance product in competition with a traditional indemnity insurance product. The index insurance product should be attractive enough to convince a sufficient number of policyholders in order to ensure the solvency of the portfolio. Our analysis suggests that the number of policyholders willing to accept index insurance can be influenced by factors such as the risk aversion of policyholders, the delay in compensation of the competing indemnity-based insurance product, and the price of the index insurance product.

We also propose the use of a hybrid product, where index insurance is used as compensation only in situations where it is most suitable for policyholders, with indemnity-based insurance used otherwise. This combination has several advantages, including accelerating compensation and reducing the premiums paid by policyholders.

Additionally, we developed an algorithm to help insurers identify the specific losses for which a compensations using index insurance are more likely to be preferred and accepted by policyholders and those for which compensations using indemnity-based insurance are more suitable. Beyond the concept of index-based insurance, the results from the hybrid insurance section can also be applied for claim management in traditional indemnity-based insurance contracts. This could be relevant in situations where the insurer aims to reduce expert costs when handling certain claims or when an insurer wishes to apply different claims management policies to various segments of their portfolio based on specific criteria.

It is important to note that we do not cover one of the most appealing aspects of index insurance in this paper, which is offering coverage for claims not covered by traditional indemnity contracts. The present paper focuses solely on the introduction of index insurance in a context

where traditional indemnity-based insurance is already in place. The question of modeling the demand for index insurance in situations where indemnity-based insurance is unavailable will be addressed in future work and will require a more nuanced approach to calibrate the utility function.

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

6 Appendix

6.1 Additional information

Coefficients	Estimate	Standard error	t value	p-value
(Intercept)	4.30287	0.25261	17.033	$< 2 \times 10^{-16}***$
T	3.10957	0.02757	112.778	$< 2 \times 10^{-16}***$
δ	-2.16544	0.11037	-19.620	$< 2 \times 10^{-16}***$
Λ	-2.17856	0.03986	-54.654	$< 2 \times 10^{-16}***$
B	-3.53310	0.32735	-10.793	$< 2 \times 10^{-16}***$

Significance codes: 0 ‘***’, 0.001 ‘**’, 0.01 ‘*’, 0.05 ‘.’, 0.1 ‘ ’

Multiple R-squared: 0.5990 *Adjusted R-squared:* 0.5987

Table 3: Results of the linear regression of Y on \mathbf{W}




Figure 8: Densities of losses due to cloud service interruptions and densities of durations of interruption. The densities are plotted for all policyholders and separately for policyholders with backup plans and those without backup plans.

6.2 Proof of Proposition 3.1

Showing the result for $\tau = 0$ is sufficient, since a higher value of τ reduces the expected utility of the traditional insurance contract.

If $\tau = 0$, condition (2.1) is

$$E [\exp(\alpha[m_Y(\alpha|\mathbf{W}) - \phi_\beta(\mathbf{W})])] \leq \exp(\alpha[1 + \theta_Y - \beta - \beta\theta]E[Y]).$$

So taking the logarithm, we get

$$\log E [\exp(\alpha[m_Y(\alpha|\mathbf{W}) - \phi_\beta(\mathbf{W})])] \leq \alpha[1 + \theta_Y - \beta - \beta\theta]E[Y].$$

Finally,

$$\theta\beta \leq 1 - \beta + \theta_Y - \frac{\log E [\exp(\alpha[m_Y(\alpha|\mathbf{W}) - \phi_\beta(\mathbf{W})])]}{\alpha E[Y]}. \quad (6.1)$$

Since $\theta\beta$ needs to be strictly positive, this requires

$$\frac{\log E [\exp(\alpha[m_Y(\alpha|\mathbf{W}) - \phi_\beta(\mathbf{W})])]}{\alpha E[Y]} < 1 - \beta + \theta_Y.$$

Since

$$\log E [\exp(\alpha[m_Y(\alpha|\mathbf{W}) - \phi_\beta(\mathbf{W})])] < \sup_{\mathbf{w} \in \mathcal{W}} \frac{m_Y(\alpha|\mathbf{w}) - \phi_\beta(\mathbf{w})}{E[Y]},$$

it follows from (3.3) that this condition holds. Then, taking $\theta \leq \eta(\alpha, \beta)\beta^{-1}$ ensures that (6.1), hence (2.1), holds.

6.3 Proof of Proposition 3.3 and 3.4

Proof of Proposition 3.3.

We have $\pi_{\phi_\beta}^* = \beta E[Y]$, and

$$\sigma_{\phi_\beta}^2 = \beta^2 \text{Var}(E[Y|\mathbf{W}]) = \beta^2 \sigma^2.$$

From Corollary 3.2, which requires that $\theta \leq \eta\beta^{-1}$,

$$n \geq n_0 = N\mu((0, \alpha_0 + h_\beta(\tau))).$$

Condition (2.3) holds if

$$\frac{n_0^{1/2} \theta \pi_{\phi_\beta}^*}{\sigma_{\phi_\beta}} \geq S^{-1}(\varepsilon).$$

This rewrites

$$\theta \geq \frac{\sigma S^{-1}(\varepsilon)}{N^{1/2} E[Y] \mu((0, \alpha_0 + h_\beta(\tau)))^{1/2}}.$$

To be compatible with the upper bound on θ , we need condition (3.6).

Proof of Proposition 3.4.

The proof is similar to the proof of Proposition 3.3, but with (2.4) replaced by (2.5). If

$$\theta \geq \frac{a\sigma(1 - \varepsilon'^\gamma)}{\gamma \varepsilon'^\gamma},$$

we have

$$S^{-1} \left(\varepsilon - \frac{1}{\left(1 + \frac{\gamma\theta}{a\sigma}\right)^{1/\gamma}} \right) \leq S^{-1}(\varepsilon - \varepsilon').$$

Again, we need the upper and lower bounds on θ to be compatible, which leads to the result of the Proposition.

6.4 Risk aversion heredity property

Lemma 6.1 *Assume that condition (2.1) holds for ϕ , $\alpha_0 > 0$ and with $\tau = 0$. Let $F(\alpha) = \Psi'(\alpha) = E[Y \exp(\alpha Y)]$. Then, (2.1) holds for all τ and $\alpha \in (0, \alpha_0 + h(\tau)]$ with*

$$h(\tau) = F^{-1}(F(\alpha_0) \exp(-\tau) \exp(-\alpha[\pi_Y - \pi_\phi])) - \alpha_0.$$

Proof. Since $\Psi(\alpha') \geq 1 = \Psi(0)$, if (2.1) holds for $\tau = 0$, it also holds for all $\tau > 0$. Hence we show the result in two steps. First, we show that condition (2.1) holds for $\tau = 0$ and all $\alpha \leq \alpha_0$. Then we study the case $\alpha_0 < \alpha \leq \alpha_0 + h(\tau)$.

First case: $\alpha \leq \alpha_0$.

Let $Z = Y - \phi(\mathbf{W}) - \pi_Y + \pi_\phi$. Condition (3.1) rewrites

$$E[\exp(\alpha Z)] \leq \Psi(\alpha').$$

We have, as a consequence of Jensen's inequality,

$$\frac{\log E[\exp(\alpha Z)]}{\alpha} \leq \frac{\log E[\exp(\alpha_0 Z)]}{\alpha_0}.$$

Hence

$$E[\exp(\alpha Z)] \leq 1.$$

Second case: $\alpha_0 < \alpha \leq \alpha_0 + h(\tau)$.

Let $\alpha = \alpha_0 + x$ for $0 < x \leq h$, and $\alpha' = \alpha(1 - \exp(-\tau))$. From a Taylor expansion,

$$\Psi(\alpha') \geq \Psi(\alpha'_0) + x(1 - \exp(-\tau))E[Y \exp(\alpha'_0 Y)].$$

On the other hand,

$$E[\exp(\alpha Z)] \leq E[\exp(\alpha_0 Z)] + x \exp(-\alpha\{\pi_Y - \pi_\phi\})E[Y \exp(\alpha Y)].$$

Since $E[\exp(\alpha_0 Z)] \leq 1 \leq \Psi(\alpha'_0)$, condition (2.1) holds if

$$\frac{F(\alpha)}{F(\alpha_0)} \leq \exp(-\tau) \exp(-\alpha[\pi_Y - \pi_\phi]).$$

By definition, this condition holds for $h \leq h(\tau)$. ■

References

- 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.
- M. Braun and A. Muermann. The impact of regret on the demand for insurance. *Journal of Risk and Insurance*, 71(4):737–767, 2004.
- L. Breiman. *Classification and regression trees*. Routledge, 2017.
- M. Carter, A. de Janvry, E. Sadoulet, and A. Sarris. Index insurance for developing country agriculture: a reassessment. *Annual Review of Resource Economics*, 9:421–438, 2017.
- L. Cesarini, R. Figueiredo, B. Monteleone, and M. L. Martina. The potential of machine learning for weather index insurance. *Natural Hazards and Earth System Sciences*, 21(8):2379–2405, 2021a.
- L. Cesarini, R. Figueiredo, B. Monteleone, and M. L. V. Martina. The potential of machine learning for weather index insurance. *Natural Hazards and Earth System Sciences*, 21(8):2379–2405, 2021b. doi: 10.5194/nhess-21-2379-2021. URL <https://nhess.copernicus.org/articles/21/2379/2021/>.
- Z. Chen, Y. Lu, J. Zhang, and W. Zhu. Managing weather risk with a neural network-based index insurance. *Management Science*, 2023.
- D. J. Clarke. A theory of rational demand for index insurance. *American Economic Journal: Microeconomics*, 8(1):283–306, 2016.
- K. Y. Clement, W. W. Botzen, R. Brouwer, and J. C. Aerts. A global review of the impact of basis risk on the functioning of and demand for index insurance. *International Journal of Disaster Risk Reduction*, 28:845–853, 2018.
- S. Conradt, R. Finger, and M. Spörri. Flexible weather index-based insurance design. *Climate Risk Management*, 10:106–117, 2015.
- J. D. Cummins and O. Mahul. The demand for insurance with an upper limit on coverage. *Journal of Risk and Insurance*, 71(2):253–264, 2004.
- L. Eeckhoudt and M. Kimball. Background risk, prudence, and the demand for insurance. *Contributions to insurance economics*, pages 239–254, 1992.
- Y. Fujii, M. Okura, and Y. Osaki. Regret, rejoicing, and mixed insurance. *Economic Modelling*, 58:126–132, 2016.

- X. Giné, R. Townsend, and J. Vickery. Statistical analysis of rainfall insurance payouts in southern india. *American Journal of Agricultural Economics*, 89(5):1248–1254, 2007.
- X. Han, G. Zhang, Y. Xie, J. Yin, H. Zhou, Y. Yang, J. Li, and W. Bai. Weather index insurance for wind energy. *Global Energy Interconnection*, 2(6):541–548, 2019.
- M. Hao, A. S. Macdonald, P. Tapadar, and R. G. Thomas. Insurance loss coverage and demand elasticities. *Insurance: Mathematics and Economics*, 79:15–25, 2018.
- L. F. Hernández-Rojas, A. L. Abrego-Perez, F. E. Lozano Martínez, C. F. Valencia-Arboleda, M. C. Diaz-Jimenez, N. Pacheco-Carvajal, and J. J. García-Cárdenas. The role of data-driven methodologies in weather index insurance. *Applied Sciences*, 13(8):4785, 2023.
- L. Johnson. Paying ex gratia: Parametric insurance after calculative devices fail. *Geoforum*, 125:120–131, 2021.
- X. Liang, W. Jiang, and Y. Zhang. Optimal insurance design under mean-variance preference with narrow framing. *Insurance: Mathematics and Economics*, 112:59–79, 2023.
- Y.-H. Lin, L.-J. Wang, X.-Y. Shi, and M.-P. Chen. Evolution of research on climate risk insurance: A bibliometric analysis from 1975 to 2022. *Advances in climate change research*, 14(4): 592–604, 2023.
- W.-Y. Loh. Fifty years of classification and regression trees. *International Statistical Review*, 82(3):329–348, 2014.
- O. Lopez and M. Thomas. Parametric insurance for extreme risks: the challenge of properly covering severe claims. *arXiv preprint arXiv:2301.07776*, 2023.
- V. R. Marone and A. Sabety. When should there be vertical choice in health insurance markets? *American Economic Review*, 112(1):304–342, 2022.
- T. Mikosch and A. V. Nagaev. Large deviations of heavy-tailed sums with applications in insurance. *Extremes*, 1(1):81–110, 1998.
- O. Prokopchuk, I. Prokopchuk, and G. Mentel. Index insurance as an innovative tool for managing weather risks in the agrarian sector of economics. *Journal of Competitiveness*, 10(3): 119, 2018.
- M. Scherer and G. Stahl. The standard formula of solvency ii: a critical discussion. *European Actuarial Journal*, 11:3–20, 2021.
- K. Tam, B. Chang, R. Hopcraft, K. Moara-Nkwe, and K. Jones. Quantifying the econometric loss of a cyber-physical attack on a seaport. *Frontiers in Computer Science*, 4:1057507, 2023.

- J. D. Woodard. Crop insurance demand more elastic than previously thought. *Choices*, 31(3), 2016.
- J. Zhang. Blended insurance scheme: A synergistic conventional-index insurance mixture. *Insurance: Mathematics and Economics*, 119:93–105, 2024.
- J. Zhang, K. S. Tan, and C. Weng. Index insurance design. *ASTIN Bulletin: The Journal of the IAA*, 49(2):491–523, 2019.