



HAL
open science

Market-based insurance ratemaking

Pierre-Olivier Goffard, Pierrick Piette, Gareth W. Peters

► **To cite this version:**

Pierre-Olivier Goffard, Pierrick Piette, Gareth W. Peters. Market-based insurance ratemaking. 2023.
hal-04297811v2

HAL Id: hal-04297811

<https://hal.science/hal-04297811v2>

Preprint submitted on 12 Sep 2024

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.

Market-based insurance ratemaking

Pierre-Olivier Goffard ^{*1}, Pierrick Piette ^{†2,3}, and Gareth W. Peters ^{‡4}

¹Université de Strasbourg, Institut de Recherche Mathématique Avancée, Strasbourg, France

²Univ Lyon, Université Claude Bernard Lyon 1, Institut de Science Financière et d'Assurances (ISFA), Laboratoire SAF EA2429, F-69366, Lyon, France

³Seyna, 10 Rue du Faubourg Montmartre 75009 Paris

⁴University of California Santa Barbara, Department of Statistics and Applied Probability, Santa Barbara CA 93106-3110, USA

September 12, 2024

Abstract

This paper introduces a method for pricing insurance policies using market data. The approach is designed for scenarios in which the insurance company seeks to enter a new market lacking historical data. The methodology involves an iterative two-step process. First, a suitable parameter is proposed to characterize the underlying risk. Second, the resulting pure premium is linked to the observed commercial premium using an isotonic regression model. To validate the method, comprehensive testing is conducted on synthetic data, followed by its application to a dataset of actual pet insurance rates. To facilitate practical implementation, we have developed an R package called `IsoPriceR`. By addressing the challenge of pricing insurance policies in the absence of historical data, this method contributes to enhancing pricing strategies in emerging markets.

MSC 2010: 62P05, 91G70, 62F15.

Keywords: Insurance Pricing, Bayesian Inference, Approximate Bayesian Computation, Isotonic Regression.

1 Introduction

Modern insurance pricing relies on predictive modeling methods to ensure that premiums reflect, as accurately as possible, the average cost of claims. To achieve this, insurers rely on historical data to train statistical

*Email: goffard@unistra.fr.

†Email: pierrick.piette@gmail.com.

‡Email: garethpeters@ucsb.edu.

learning models and calculate what is called the pure premium. While the foundation of standard actuarial practice often rests on generalized linear models (GLM), see Renshaw [14], the relentless evolution of data science has ushered in a new era where more sophisticated machine learning algorithms are also coming into play, see Blier-Wong et al. [4] and the reference therein. However, a challenge arises when an insurance company enters a new market, lacking historical data on the risks it aims to cover. In this context, conventional predictive modeling tools are failing, leaving insurers at a crossroads, looking for innovative solutions to navigate uncharted territory.

Although an insurance company may lack historical data in a new market, there is an attractive alternative: the ability to observe and analyze market data consisting of rates offered by competitors for similar insurance policies. Our approach leverages this market data to provide insights into the underwritten risk leading to the calculation of insurance premiums. [Our objective is to develop a methodology that determines suitable commercial premiums based on the observed commercial rates of competitors.](#)

[The data collection process involves obtaining insurance quotes. To gather these quotes, one typically visits insurance company websites and answers several questions about the insured risk. The premiums quoted depend on the responses provided by the customer. In this paper, we use pet insurance as a case study. Pet insurance covers health expenses for pets. A pet owner is required to specify the pet's specie, breed, age, and gender. These characteristics are referred to as rating factors in actuarial science and are crucial for risk classification. For an overview of this topic, we direct the reader to the work of Antonio and Valdez \[1\]. An insurance company looking to enter a new market would naturally identify these risk factors when collecting data on the premiums offered by competitors.](#)

[Within a specific risk class, the risk is represented by a positive random variable, denoted as \$X\$, which quantifies the total claim amount over the insurance policy period. Insurance companies mitigate this risk by offering coverage for a portion of \$X\$, denoted as \$g\(X\) < X\$, in exchange for a premium. The process involves calculating the pure premium, defined as \$p = \mathbb{E}\[g\(X\)\]\$. Customers are subsequently offered a commercial premium that is related to the pure premium by the equation \$\pi = f\(p\) > p\$, where \$f\$ represents the loading function. Our problem is formulated as follows: given a collection of insurance quotes \$\mathcal{D} = \{\pi_1, \dots, \pi_n\}\$ associated with a specific risk and varying loading and insurance coverage functions, represented as](#)

$$\pi_i = f_i\{\mathbb{E}[g_i(X)]\}, \quad i = 1, \dots, n,$$

[we aim at studying the distribution of the underlying risk \$X\$ as well as approximate the loading functions in order to quote our own insurance policies.](#)

Our approach assumes that the distribution of the risk X is parameterized by $\theta \in \mathbb{R}^d$. Ideally, if the loading functions f_i , were known for $i = 1, \dots, n$, a procedure akin to the generalized method of moments could be used, see Hansen [10]. Unfortunately, these functions are not known. We adopt a two steps procedure. We sample θ from a prior distribution $p(\theta)$ which delimits our parameter space. We compute the pure premiums $p_i^\theta = \mathbb{E}_\theta[g_i(X)]$, based on the known coverages g_i provided by the insurance policy and the sampled parameter value θ . One key challenge to address is the lack of tractable expressions for the pure premiums as we assume that X is governed by a compound distribution (i.e. a random sum). We therefore use a crude Monte-Carlo approach which entails an approximation error linked to the number of Monte-Carlo replications of X we use. The loading functions $f_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ are approximated using an isotonic regression model, chosen for its ability to maintain the monotonic relationship between pure and commercial premiums—a desirable feature. Additionally, market data is inherently noisy, and isotonic regression provides robustness to outliers, superior to that of simple linear regression, for example. The procedure may be summarized as follows:

1. Sample a parameter value θ
2. Compute the pure premiums $p_i^{\theta^*}$ for each of the insurance policies $i = 1, \dots, n$
3. Fit an isotonic function f to learn the relationship between the commercial premia π_i and the pure premia p_i^θ
4. Build the 'synthetic' market data \mathcal{D}^θ by applying the estimated loading f to the pure premium $f(p_i^\theta)$ for $i = 1, \dots, n$
5. If the observed and synthetic market data are close enough then we store the parameter value θ and the associated loading function f .

After iterating the above steps, we get a sequence of parameter-loading function pairs: $(\theta_1, f_1), (\theta_2, f_2), \dots$. This sequence allows us to price our own insurance policies. The problem we tackle is an inverse problem and our solution is inspired from indirect inference methodologies pioneered by Gourieroux et al. [9]. The proposed algorithm to search the parameter space resembles Approximate Bayesian Computation (ABC) algorithms described in the book of Sisson et al. [16]. The parameter values sampled $\theta_1, \theta_2, \dots$ by the algorithm yields an approximate posterior distribution $p(\theta|\mathcal{D})$. This posterior distribution accounts for the uncertainty around the estimated parameter value due to the use of Monte Carlo simulation to calculate the pure premium. ABC algorithms have found successful applications in a range of actuarial science and risk management problems.

We refer the readers to the works of Peters et al. [13], Dean et al. [5], Peters and Sisson [12] and Goffard and Laub [8] for further insights. Isotonic regression, a well-established statistical methodology, see for instance Barlow et al. [2], plays a central role in our approach. A recent application in actuarial science addresses the autocalibration challenges that can arise when pricing insurance contracts using machine learning algorithms, see the work of Wüthrich and Ziegel [17].

The rest of the paper is organized as follows. [Section 2](#) describes the risk model used in this study and discusses insurance pricing principles. [Section 3](#) provides a detailed account of the algorithmic procedure. Our method is presented as an Approximate Bayesian Computation (ABC)-type optimization algorithm, which incorporates a simple isotonic regression model. [Section 4](#) presents the results of a simulation study designed to showcase the performance of our method in a controlled environment. Lastly, we apply our algorithm on a dataset made of real-world pet insurance rates in [Section 5](#).

2 Model set up and insurance premium computation

An individual seeks to hedge against a risk X modeled by a positive random variable, over a given period of time, say one year. A common model used for X in property and casualty insurance is given by a compound loss variable

$$X = \sum_{k=1}^N U_k, \quad (1)$$

where N is a counting random variable and the U_k 's are independent and identically distributed (i.i.d) positive random variables independent from N . The random variable N is the number of occurrences of an event over a given time period (annually), each of these events is associated to a compensation U_k . [We assume here that \$X\$ represents a risk that belongs a specific risk class determined by risk factors.](#)

2.1 Pure premium computation

An insurance company offers to bear part of this risk $g(X) \leq X$ in exchange for a premium which should compensate the average cost of claim given by $p = \mathbb{E}[g(X)]$, referred to as the pure premium. We consider in this work a function g defined as

$$g(x) = \min(\max(r \cdot x - d, 0), l),$$

where $r \in (0, 1]$ is the coverage rate, $d > 0$ is the deductible and $l > 0$ is the limit. [We illustrate the impact of the parameters of the insurance coverage in Example 1.](#)

Example 1. Let us consider a scenario where the risk has a Poisson-lognormal distribution $X \sim \text{Poisson}(\lambda = 3) - \text{LogNorm}(\mu = 0, \sigma = 1)$ and that $n = 100$ insurance coverages are proposed. These are characterized by a rate, a deductible and a limit, set randomly as

$$r_i \sim \text{Unif}([0.5, 1]), d_i \sim \text{Unif}([0.5, 6]), \text{ and } l = \infty, \text{ for } i = 1, \dots, 100.$$

Figure 1 shows the pure premiums

$$p_i = \mathbb{E}(g_i(X)) = \mathbb{E}(\min(\max(r_i \cdot x - d_i, 0), l_i)), \quad i = 1, \dots, 100. \quad (2)$$

as a function of the rates and deductibles.




Figure 1: Pure premiums as a function of the rate of coverage (r) and the deductible (d) for a $\text{Poisson}(\lambda = 3) - \text{LogNorm}(\mu = 0, \sigma = 1)$ risk.

The pure premium are increasing in the rates and decreasing in the deductible. Note that the pure premiums were estimated via a Crude Monte Carlo simulation method to overcome the lack of explicit formula for the distribution functions of X .

In practice, the rate offered to policyholders include a loading to compensate for the variability of the risk and cover the management costs. We describe this loading in the next section.

2.2 From pure premiums to commercial premiums

Let $f : \mathbb{R}_+ \mapsto \mathbb{R}_+$ be a nondecreasing function, such that

$$\pi = f(p) \geq p.$$

The function f is referred to as the loading function. As the commercial premium is a function of the pure premium then we are applying the expectation premium principle. Other premium principle are also possible like the standard deviation principle discussed in [Appendix B](#). A simple loading function is linear in the pure premium as

$$f(x) = (1 + \eta)x,$$

where $\eta > 0$. The loading functions used by insurance companies are unknown to us and vary from one insurance company to the other. [We follow up on Example 1 in Example 2 where we randomize the linear link between pure and commercial premium.](#)

Example 2. Take the pure premiums of [Example 1](#) and apply the following linear loadings

$$\eta_i \sim \text{Unif}([0.5, 2]), \text{ for } i = 1, \dots, n.$$

The commercial premium then relates to the pure premium as

$$\pi_i = (1 + \eta_i)p_i, \text{ for } i = 1, \dots, n. \tag{3}$$

[Figure 2](#) displays the commercial premium as a function of the pure premium.




Figure 2: Pure premium as a function of the commercial premium offered by various insurance companies.

We only observe the commercial premium π_1, \dots, π_n and we would like to learn from them about the risk X and the loading functions f_1, \dots, f_n . We formulate our problem and describe our solution in the next section.

3 Market derived insurance ratemaking

The data at hand is a collection of insurance rates $\mathcal{D} = \{\pi_1, \dots, \pi_n\}$ associated to insurance coverages g_1, \dots, g_n for a given risk X within a particular risk class. We suppose that these insurance rates \mathcal{D} were derived according to the model described in Assumption 1.

Assumption 1. *The commercial rate $\mathcal{D} = \{\pi_1, \dots, \pi_n\}$ are given by*

$$\pi_i = f_i(p_i) = f_i\left\{\mathbb{E}_{\theta_0}[g_i(X)]\right\}, \quad i = 1, \dots, n,$$

where the risk X is a random variable defined in (1), and parametrized by an unknown parameter $\theta_0 \in \Theta \subset \mathbb{R}^d$. The loading functions f_1, \dots, f_n are unknown. The insurance coverage functions g_1, \dots, g_n are known and of the form

$$g_i(x) = \min(\max(r_i \cdot x - d_i, 0), l_i), \quad \text{for } i = 1, \dots, n.$$

We wish to identify the parameter θ that best explains our data based on a loading function f that attempts to average out the loading functions f_1, \dots, f_n used by the competitors. We aim to solve the following optimization problem.

Problem 1. *Denote by*

$$p_i^\theta = \mathbb{E}_\theta[g_i(X)], \quad \text{for } i = 1, \dots, n,$$

the pure premium associated to a risk X parametrized by $\theta \in \Theta \subset \mathbb{R}^d$. Further denote by $p_{1:n}^\theta$ and $\pi_{1:n}$ the collections of pure and commercial premiums.

We wish to find $\theta \in \Theta \subset \mathbb{R}^d$ and $f : \mathbb{R}_+ \mapsto \mathbb{R}_+$ to minimize

$$d\left[\pi_{1:n}, f(p_{1:n}^\theta)\right], \quad (4)$$

where the function f is applied elementwise on $p_{1:n}^\theta$ and $d(\cdot, \cdot)$ denotes a distance function over the observation space, subject to

$$\pi_i \geq p_i^\theta \quad (5)$$

and

$$f(p_i^\theta) \geq p_i^\theta, \quad (6)$$

for $i = 1, \dots, n$.

The link between pure and commercial premiums is approximated using an isotonic regression model described in Section 3.1. Section 3.2 focuses on the distance (4) used to compare the observed and model-generated commercial rates. Finally, Section 3.3 refines the accept-reject algorithm laid out in the introduction to search the parameter space in a more efficient way.

3.1 Isotonic regression

Isotonic regression is a statistical technique used for fitting a non-decreasing function to a set of data points. Suppose that we hold a collection of candidate pure premiums $p_{i:n}^\theta$, associated to a candidate risk parameter θ . Our datapoints are therefore pairs of pure and commercial premiums $(p_i^\theta, \pi_i)_{i=1, \dots, n}$. Suppose the pure premium have been ordered such that $p_i^\theta \leq p_j^\theta$ for $i \leq j$, isotonic regression seeks a least square fit π_i^θ for the π_i 's such that $\pi_i^\theta \leq \pi_j^\theta$ for $p_i \leq p_j$. It reduces to find $\pi_1^\theta, \dots, \pi_n^\theta$ that minimize

$$\sum_{i=1}^n (\pi_i^\theta - \pi_i)^2, \text{ subject to } \pi_i^\theta \leq \pi_j^\theta \text{ whenever } p_i^\theta \leq p_j^\theta.$$

Since the p_i 's fall in a totally ordered space, a simple iterative procedure called the Pool Adjacent Violators Algorithm (PAVA) can be used. Here's a high-level overview of how it works:

1. Initialize the sequence of values to be the same as the data points $\pi_i^* = \pi_i$.
2. Iterate through the sequence and identify "violations," which occur when the current value is greater than the next value, that is

$$\pi_i^* > \pi_{i+1}^* \text{ for some } i = 1, \dots, n.$$

When a violation is found, adjust the values in the associated segment of the sequence to be the average of the values,

$$\pi_i^* \leftarrow (\pi_i^* + \pi_{i+1}^*)/2,$$

ensuring monotonicity.

3. Repeat Step 2 until no violations are left.

We use the `isoreg` function from *R* to get the fitted values π_i^θ , $i = 1, \dots, n$. To complete the isotonic regression task we shall find a function f such that $f(p_i^\theta) = \pi_i^\theta$. A common choice is a piece-wise constant function that interpolates the π_i^θ 's as illustrated in [Example 3](#).

Example 3. *The isotonic fit of the data of [Example 1](#) and [Example 2](#) is provided on [Figure 3](#).*

Remark 3.1. *When looking at [Figure 3](#), one may object that a simple linear regression model could do the job. This impression is partly due to the (noisy) linear link between pure and commercial premium in (14). Isotonic regression is a non-parametric approach, meaning it doesn't make strong assumptions about the underlying distribution or functional form of the relationship between variables. This can be advantageous when the true relationship is not well represented by a linear model. Furthermore, Isotonic regression is generally more resistant to outliers compared to*




Figure 3: Isotonic link between the pure and commercial premiums.

linear regression due to its piecewise constant nature and the way it enforces monotonicity. We briefly illustrate this fact in Appendix A by looking at the residuals of the linear and isotonic regression and considering a non-linear link function between pure and commercial premium.

The choice of isotonic regression anticipates the necessity of satisfying the constraints (5) and (6) of Problem 1. This issue, which is primarily handled through the definition of an appropriate distance between observed and model-generated commercial premiums, is at the center of the next subsection.

3.2 Distance and regularization

Our starting point to measure the discrepancy between observed and model-generated commercial rates is the root mean square error (RMSE) defined as

$$\text{RMSE}[\pi_{1:n}, f(p_{i:n}^\theta)] = \sqrt{\frac{1}{n} \sum_{i=1}^n [\pi_i - f(p_i^\theta)]^2}, \quad (7)$$

for a candidate risk parameter θ and an isotonic fit f . The way isotonic regression works implies that if two parameters θ_1 and θ_2 yield the same ordering of pure premiums, then these two parameters will lead to the same RMSE. When multiple values of the parameters are equally good, the model is said to be non-identifiable. If we had access to the pure premium $p_{1:n}$ instead of the commercial rates, then our problem would reduce to finding $\theta \in \Theta$ that minimizes

$$\text{RMSE}(p_{1:n}, p_{1:n}^\theta) = \sqrt{\frac{1}{n} \sum_{i=1}^n (p_i - p_i^\theta)^2}. \quad (8)$$

In that case, the statistical framework would be that of minimum distance estimation. As we do not have access to the full shape of the data distribution, we base our inference on specific moments, just as in the generalized method of moments, a popular method among econometricians (see Hansen [10]). The consistency of the estimators depends on the functional g_i 's and is not easy to verify from a theoretical standpoint. Given the shape of our insurance coverage function, there is hope in narrowing it down to the right parameters if the pure premiums were accessible, as illustrated in Example 4.

Example 4. We consider the same model as in Example 1. Recall that the risk has a Poisson-lognormal distribution $X \sim \text{Poisson}(\lambda = 3) - \text{LogNorm}(\mu = 0, \sigma = 1)$ and that $n = 100$ insurance coverages are proposed. The rates, deductibles and limit, are set randomly as

$$r_i \sim \text{Unif}([0.5, 1]), d_i \sim \text{Unif}([0.5, 6]), \text{ and } l = \infty, \text{ for } i = 1, \dots, 100.$$

Let us assume that $\mu = 0$ and compute the pure premium over a grid of values for λ and σ . Figure 4 shows the $\text{RMSE}[p_{1:n}, p_{1:n}^\theta]$ depending on the value of λ and σ for $(\lambda, \sigma) \in [0, 5] \times [0, 2]$. This contour plot shows minimal RMSE




Figure 4: Contour plot of $\text{RMSE}(p_{1:n}, p_{1:n}^\theta)$ for $\mu = 0$ and $(\lambda, \sigma) \in [0, 5] \times [0, 2]$.

values around the true parameter values $(\lambda_0, \sigma_0) = (3, 1)$.

When dealing instead with commercial premiums and approximating the safety loadings function via an isotonic regression model, many sets of parameters could exactly reproduce the ordering of the pure premiums of the true parameter. Our problem is an ill-posed inverse problem. Ill-posedness is usually dealt with by adding a regularization to the objective function that one wants to minimize. Our solution is based on targeting a given loss ratio. The ratio of p/π corresponds to what practitioners would call the expected Loss Ratio (LR). The loss ratio is a standard measure to assess the profitability of insurance lines of business. An insurance company that enters a new market is likely to have insights on the loss ratio relative to this market , for example

by having informal discussions with reinsurers, brokers or competitors. These insights may translate into the definition of a lower and upper bound denoted by LR_{low} and LR_{high} , respectively. We can then assume that the loss ratios $\text{LR}_i = p_i/\pi_i$, for $i = 1, \dots, n$, should fall in the range $[\text{LR}_{\text{low}}, \text{LR}_{\text{high}}]$, which we refer to as the loss ratio corridor. Assuming that $\text{LR}_{\text{high}} < 1$, we may ensure both constraint (5) and $\text{LR}_i \in [\text{LR}_{\text{low}}, \text{LR}_{\text{high}}]$ by adding to our distance (7) two regularization terms defined as

$$\text{Reg}_{\text{low}}(\pi_{1:n}, p_{1:n}^\theta) = \sqrt{\frac{1}{n} \sum_{i=1}^n (\pi_i - p_i^\theta \cdot \text{LR}_{\text{low}}^{-1})_+^2},$$

and

$$\text{Reg}_{\text{high}}(\pi_{1:n}, p_{1:n}^\theta) = \sqrt{\frac{1}{n} \sum_{i=1}^n (p_i^\theta \cdot \text{LR}_{\text{high}}^{-1} - \pi_i)_+^2},$$

where $(x)_+ = \max(x, 0)$ denotes the positive part of x . The distance we consider within Problem 1 is now given by

$$d[\pi_{1:n}, f(p_{1:n}^\theta)] = \text{RMSE}[\pi_{1:n}, f(p_{1:n}^\theta)] + \text{Reg}_{\text{low}}(\pi_{1:n}, p_{1:n}^\theta) + \text{Reg}_{\text{high}}(\pi_{1:n}, p_{1:n}^\theta).$$

We illustrate the impact of adding the regularization terms in Example 5.

Example 5. We consider the commercial rates of Example 2. Recall that the commercial premiums are given by

$$\pi_i = (1 + \eta_i)p_i, \quad \text{for } i = 1, \dots, n, \quad (9)$$

where the pure premiums are those of Example 4 and

$$\eta_i \sim \text{Unif}([0.5, 2]), \quad \text{for } i = 1, \dots, n.$$

Figure 5 displays the contour plot of the discrepancy between observed and model-generated commercial rates. When comparing Figure 5a and Figure 5b, we note how beneficial including regularization terms is to identify the true parameter values.

Now that we have defined a distance function to compare observed and model-generated commercial premiums, we can describe the procedure to search the parameter space.

3.3 Population Monte Carlo Approximate Bayesian Computation algorithm

Our solution alternates between proposing parameter values for the risk to compute the pure premiums and approximating the f_i 's using isotonic regression. We must accommodate the lack of tractable expressions for the pure premium

$$p_i^\theta = \mathbb{E}_\theta[g_i(X)], \quad \text{for } i = 1, \dots, n.$$

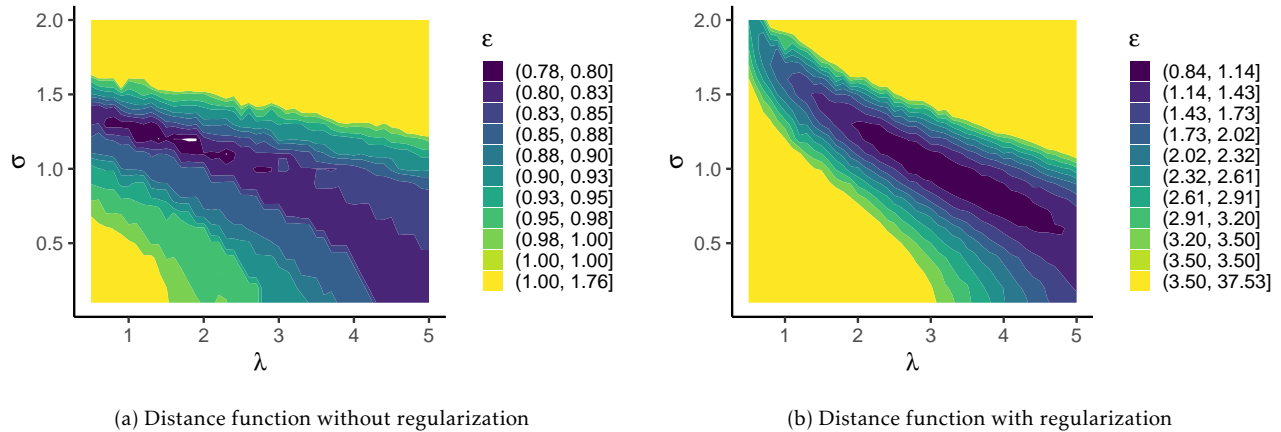


Figure 5: Contour plot of $\text{RMSE}[\pi_{1:n}, f(p_{i:n}^\theta)]$ and $d[\pi_{1:n}, f(p_{i:n}^\theta)]$ for $\mu = 0$ and $(\lambda, \sigma) \in [0, 5] \times [0, 2]$.

The use of numerical methods makes a grid search procedure prohibitive from a computing time point of view. It also prevents us from using gradient-based optimization procedures. In such cases, one can turn towards particle swarm optimization algorithms or genetic algorithms to search the parameter space. Since we have decided to take a Crude Monte Carlo estimator for the pure premiums, the accuracy depends on the number of replications R of X being used. We adopt a Bayesian strategy in order to reflect the uncertainty around the pure premium calculation onto the parameters' final estimates. Our algorithm is similar to Approximate Bayesian Computation algorithms and we simply refine the procedure laid out in the introduction to get an approximation of the posterior distribution $p(\theta|\mathcal{D})$.

We start by setting a prior distribution $p(\theta)$ over the parameter space that we sequentially improve through intermediate distributions characterized by a sequence of tolerance levels $(\epsilon_g)_{g \geq 0}$ that decrease gradually as $\infty = \epsilon_0 > \epsilon_1 > \epsilon_2 > \dots > 0$. Each intermediate distribution (called a generation) is represented by a cloud of weighted particles $(\theta_j, w_j)_{j=1, \dots, J}$. We approximate each intermediate posterior distribution using a multivariate kernel density estimator (κDE) denoted by $p_{\epsilon_g}(\theta|\mathcal{D})$. The parameters of the algorithm are the number of generations G and the population size J (the number of particles in the cloud).

The algorithm is initialized by setting $\epsilon_0 = \infty$ and $p_{\epsilon_0}(\theta|\mathcal{D}) = p(\theta)$. For generation $g \geq 1$, we hold an intermediate distribution $p_{\epsilon_{g-1}}(\theta|\mathcal{D})$ from which we can sample particles $\theta^* \sim p_{\epsilon_{g-1}}(\theta|\mathcal{D})$. We compute the associated pure premium

$$p_i^{\theta^*} = \mathbb{E}_{\theta^*}[g_i(X)], \quad \text{for } i = 1, \dots, n.$$

The pure premiums are computed via Monte Carlo simulation. The accuracy depends on the number R of

copies of X involved in the Monte Carlo estimations. We then fit the isotonic regression model

$$\pi_i = f(p_i^{\theta^*}) + e_i, \quad \text{for } i = 1, \dots, n,$$

where e_i is an error term that captures the mismatch between the true value of the pure premium and its empirical counterpart estimated by the competitor insurance company using its historical data and the company-specific loading function. We further compare the observed commercial premiums to the model-generated ones via the distance defined in [Section 3.2](#) with

$$d[\pi_{1:n}, f(p_{1:n}^{\theta^*})] = \text{RMSE}[\pi_{1:n}, f(p_{1:n}^{\theta^*})] + \text{Reg}_1(\pi_{1:n}, p_{1:n}^{\theta^*}) + \text{Reg}_2(\pi_{1:n}, p_{1:n}^{\theta^*}).$$

If the distance satisfies $d[\pi_{1:n}, f(p_{1:n}^{\theta^*})] < \epsilon_{g-1}$, then we keep the associated particle θ^* . New particles are proposed until we reach J accepted particles denoted by $\theta_1^g, \dots, \theta_J^g$. We also store the distances d_1^g, \dots, d_J^g . We need to set the next tolerance threshold ϵ_g , which is used to calculate the particle weights

$$w_j^g \propto \frac{p(\theta_j^g)}{p_{\epsilon_{g-1}}(\theta)} \mathbb{I}_{d_j^g < \epsilon_{g-1}}, \quad j = 1, \dots, J.$$

The tolerance threshold is chosen so as to maintain a specified effective sample size (ESS) of $J/2$ as in Del Moral et al. [6]. Following Kong et al. [11], the ESS is estimated by $1/\sum_{j=1}^J (w_j^g)^2$. This weighted sample then allows us to update the intermediate distribution as

$$p_{\epsilon_g}(\theta|\mathcal{D}) = \sum_{j=1}^J w_j^g K_H(\theta - \theta_j^g),$$

where K_H is a multivariate KDE with smoothing matrix H . A common choice for the KDE is the multivariate Gaussian kernel with a smoothing matrix set to twice the empirical covariance matrix of the cloud of particles $\{\theta_j^g, w_j^g\}$ as in Beaumont et al. [3]. The procedure is summarized in [Algorithm 1](#).

The user must configure several aspects of the algorithm. The prior assumptions $p(\theta)$ determine the parameter space that will be searched. The loss ratio corridor $[\text{LR}_{\text{low}}, \text{LR}_{\text{high}}]$ sets up the two regularization terms, ensuring that parameters associated with unreasonable pure premiums are excluded. Prior settings and the loss ratio corridor can be guided by expert opinions. The population size J drives the quality of the posterior distributions approximations through the cloud of particles. A large J also enhances the chances of finding the global optimum, as more particles improve the coverage of the parameter space. A greater number R of Monte Carlo simulations ensures the accuracy of the pure premium evaluation. Both R and J contribute to the stability of the algorithm's results over several runs. The number of generations G relates to the tolerance level ϵ , which in turn drives the narrowness of the posterior distribution output by the ABC algorithm. As one

Algorithm 1 Population Monte Carlo Approximate Bayesian Computation

- 1: **set** $\epsilon_0 = \infty$ and $p_{\epsilon_0}(\theta | \mathcal{D}) = \pi(\theta)$
 - 2: **for** $g = 1 \rightarrow G$ **do**
 - 3: **for** $j = 1 \rightarrow J$ **do**
 - 4: **repeat**
 - 5: **generate** $\theta^* \sim p_{\epsilon_{g-1}}(\theta | \mathbf{x})$
 - 6: **compute** $p_i^{\theta^*} = \mathbb{E}_{\theta^*}[g_i(X)]$, for $i = 1, \dots, n$
 - 7: **fit** the isotonic regression model $\pi_i = f(p_i^{\theta^*}) + e_i$, for $i = 1, \dots, n$
 - 8: **compute** $d[\pi_{1:n}, f(p_{1:n}^{\theta^*})] = \text{RMSE}[\pi_{1:n}, f(p_{1:n}^{\theta^*})] + \text{Reg}_1(\pi_{1:n}, p_{1:n}^{\theta^*}) + \text{Reg}_2(\pi_{1:n}, p_{1:n}^{\theta^*})$.
 - 9: **until** $d[\pi_{1:n}, f(p_{1:n}^{\theta^*})] < \epsilon_g$
 - 10: **set** $\theta_j^g = \theta^*$ and $d_j^g = d^*$
 - 11: **end for**
 - 12: **find** $\epsilon_g \leq \epsilon_{g-1}$ so that $\widehat{\text{ESS}} = \left[\sum_{j=1}^J (w_j^g)^2 \right]^{-1} \approx J/2$, where
$$w_j^g \propto \frac{p(\theta_j^g)}{p_{\epsilon_{g-1}}(\theta_j^g | \mathcal{D})} \mathbb{I}_{d_j < \epsilon_g}, \quad j = 1, \dots, J$$
 - 13: **compute** $p_{\epsilon_g}(\theta | \mathcal{D}) = \sum_{j=1}^J w_j^g K_H(\theta - \theta_j^g)$
 - 14: **end for**
-

would expect, the computational time for the algorithm increases with higher values of R, J and G . Therefore, the choice of suitable values for G, J , and R can be made in consideration of a predetermined computational time budget. A practical solution to set G on the fly is to stop the algorithm whenever the difference between two consecutive tolerance levels is lower than some threshold Δ_ϵ . We illustrate the posterior distribution evolution along the algorithm iterations in [Example 6](#).

Example 6. We follow up on [Example 4](#) and [Example 5](#). We aim to fit the model

$$X \sim \text{Poisson}(\lambda) - \text{LogNorm}(\mu = 0, \sigma).$$

The prior assumptions are as follows

$$\lambda \sim \text{Unif}([0, 10]), \text{ and } \sigma \sim \text{Unif}([0, 5]).$$

The algorithm parameters are set to

$$J = 1000, R = 1000, \Delta_\epsilon = 0.1, \text{ and } LR \in [0.3, 0.66].$$

The algorithm halts at the 9th generation reaching a tolerance level of $\epsilon = 0.87$. [Figure 6](#) shows the sequence intermediate posterior distributions for λ and σ .

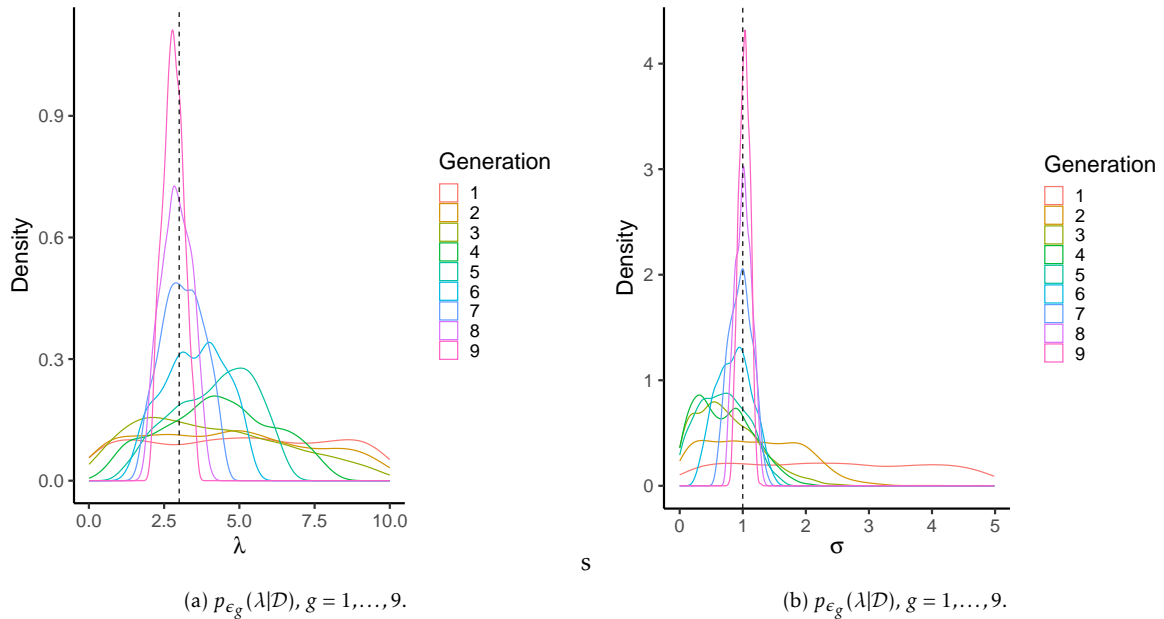


Figure 6: Intermediate posterior distributions of λ and σ .

After the algorithm terminates, it is customary to focus on the last generations of particles for inference. Pointwise estimators are derived from this final set of particles. Two commonly used estimators include the Mean *A Posteriori* (MAP) obtained by averaging the particles in the last cloud and the Mode *A Posteriori* (MODE), which is the mode of the empirical distribution within the final cloud of particles. The simulation study, conducted in the following section, is designed to investigate the convergence behavior and to compare the characteristics of the MAP and MODE estimators.

4 Methodology Assessment via Simulation

In this section, we embark on an empirical exploration, seeking to understand how the posterior distribution of the parameters behaves as the sample size n increases. This experimentation has been designed to resemble as much as possible the real data situation considered in [Section 5.3](#). We consider the risk, [within a particular risk class](#), to be distributed as the random variable

$$X = \sum_{k=1}^N U_k,$$

where

$$N \sim \text{Poisson}(\lambda = 0.3), \tag{10}$$

and

$$U_k \sim \text{LogNorm}(\mu = 6, \sigma = 1), k = 1, \dots, N. \tag{11}$$

The U_i 's are iid and independent from N . We suppose that we know the variance parameter σ and we try to draw inference on λ and μ . The parameter values of the claim frequency and severity in (10) and (11) respectively are those inferred in [Section 5.3](#) for the Poisson – LogNorm model using the MODE estimator. The prior distributions are set to independent uniforms for λ and μ as

$$\lambda \sim \text{Unif}([0, 10]), \text{ and } \mu \sim \text{Unif}([-10, 10]).$$

We generate artificial synthetic commercial premiums for this case study according to

$$\pi_i = (1 + \eta_i)\mathbb{E}[g_i(X)] = (1 + \eta_i)\mathbb{E}\{\min[\max(r_i \cdot X - d_i, 0), l_i]\}, i = 1, \dots, n,$$

where the premium parameters r , d and l are sampled from that of the real data considered in [Section 5](#), so that the simulated data is as close as possible to the real data. The η_i 's are iid from $\eta_i \sim \text{Unif}([1.43, 2.5])$, which corresponds to loss ratios between 40% and 70%. We further set $\text{LR}_{\text{low}} = 40\%$ and $\text{LR}_{\text{high}} = 70\%$. We consider sample of sizes 25, 50, 100, and 200. We configure the algorithm with a population size of $J = 1,000$ and use

$R = 2,000$ Monte Carlo replications. To ensure the algorithm’s efficiency, we set a stopping threshold, requiring that the difference between two consecutive tolerance levels is smaller than $\Delta_\epsilon = 1$ for the algorithm to halt. These settings are kept for the analysis of real-world data, as they strike a balanced compromise between accuracy and computing time. We generate 100 samples of fake data and apply our procedure. Our goal is to compare the result obtained using our two pointwise estimators: the mean *a posteriori* MAP and the mode *a posteriori* MODE. The estimators of the parameters λ and μ are given on Figure 7.

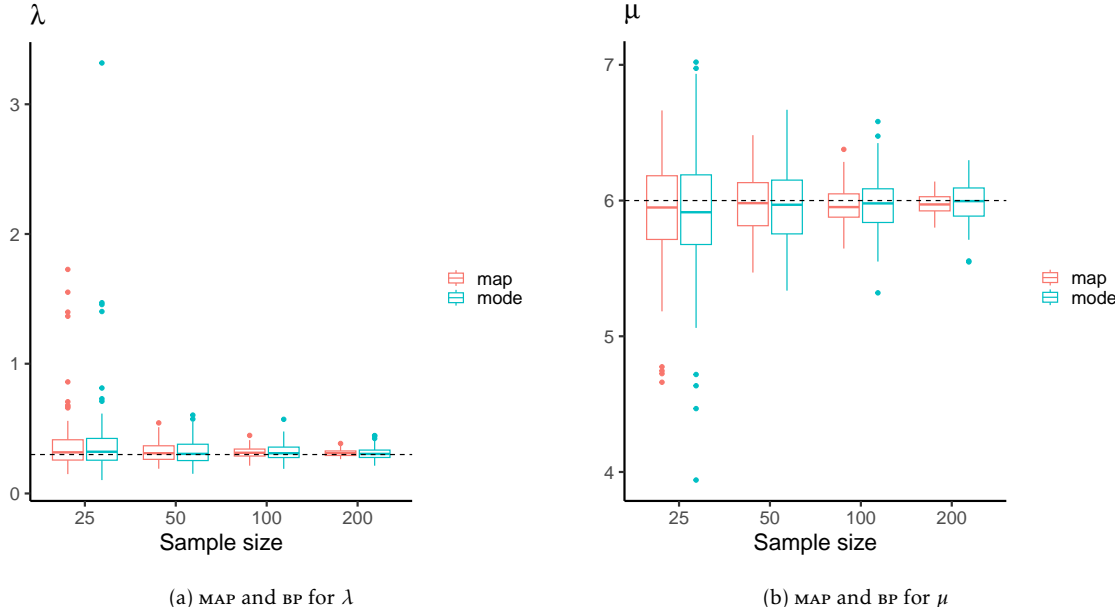


Figure 7: MAP and MODE estimators of the parameter of the model $\text{Poisson}(\lambda = 0.58) - \text{LogNorm}(\mu = 5.75, \sigma = 1)$ based on synthetic market data of sizes 50, 100, 250, and 500.

Both of the point-wise estimators seem to converge toward the parameter values that generated the data. The MAP exhibits a better behavior than the MODE as its variability decreases in a notable way as the sample size increases.

In Figure 8, we present a comparison of key metrics, including the average claim amount, the average claim frequency, the probability of no reported claims, the average total claim amount, and the average loss ratio, defined as

$$\overline{\text{LR}} = \frac{1}{n} \sum_{i=1}^n \frac{p_i}{\pi_i},$$





Figure 8: MAP and MODE estimator of the features of the Poisson($\lambda = 0.3$) – LogNorm($\mu = 6, \sigma = 1$) loss model based on synthetic market data of sizes 25, 50, 100, and 200.

Both estimation methods yield satisfactory results in recovering the characteristics of the loss distribution but the use of the MAP yields more reliable estimations.

5 Application to the pet insurance market

5.1 Evolution and growth of the pet insurance market

Pet insurance is a product designed to cover the costs of veterinary care for pets. It operates on a similar principle to human health insurance, providing a way for pet owners to manage the financial risks associated with unexpected medical expenses for their animals. Usually the expenses are covered in case of an accident

or a disease. Pet owners can choose from different policy options based on their budget and coverage needs. Policies may vary in terms of deductibles (d), coverage limits (l), and coverage rates (r). The cost of premiums can depend on various factors, including the pet's age, breed, health condition, and the level of coverage selected.

The pet insurance market has been witnessing significant growth globally, driven by increasing pet ownership (especially with so-called pandemic pets, i.e. animals adopted during 2020 lockdowns), rising veterinary costs and the changing role that a pet plays in a families social structure. This latter factor is also influenced by changing societal views and the increased awareness of the importance of health and welfare of pets, which in turn comes with increased consideration of regular veterinary health checks . In order to offset the cost associated with such expenditures, there has begun to be a broader interest in households purchasing pet insurance.

To date, the adoption and acceptance of pet insurance still varies significantly across regions of the world. Nordic countries, such as Sweden, have historically had a very high penetration rate with around 70% of pets insured. Some Anglo-Saxon countries (UK and Germany mostly) have seen significant growth in the pet insurance market during the last decades, leading to 30% of penetration rate. Other developed countries, like France, have significantly lower market sizes, with less than 10% of pets that are insured, which suggests high growth potential. The market place for pet insurance in the USA is currently also experiencing sustained growth. According to, MarketWatch guides annual insurance surveys¹, about 44.6% of pet owners stated they currently have pet insurance in the nationwide survey. Furthermore, the North American Pet Health Insurance Association (NAPHIA) undertook a survey in 2022 on the "State of the Industry Report" and found that more than 4.41 million pets were insured in North America in 2021, up from 3.45 million in 2020. The report also found that \$2.84 billion of pet insurance premiums were in force in 2021, a 30.5% increase from 2020.

This growth continues to spur increases in the capital investments associated with such an insurance line of business:

- in Sweden, Lassie has raised 11m euros in 2022 [and 23m euros in 2023](#) ;
- in the UK, ManyPets has raised \$350m at a valuation higher than \$2bn in 2021;
- in France, Dalma has raised 15m euros in 2022.
- [JAB Holding Company has invested around 2 billion dollars in 2021 to create the Pinnacle Pet Group](#)

¹<https://www.marketwatch.com/guides/pet-insurance/pet-insurance-facts-and-statistics/>

and the Independence Pet Holdings. Their purpose is to become the pet insurance leaders respectively in Europe and North America through multiple acquisitions of historic players.

Hence, the pet insurance market is becoming more competitive with an increasing number of insurance companies or brokers offering pet insurance policies. To gain new market shares as a new agent, there is a need to propose differentiated products such as new cover mixes without deductible and higher limits.

5.2 Data description

During the week of the 18th of May 2024, we have collected 1,080 quotes from 5 insurance companies. Each row of our datasets corresponds to a yearly premium collected from some insurance company website associated to a specific insurance coverage and a specific dog. We therefore find the coverage parameters which are the coverage rate r , the deductible d and the limit l . Recall that the compensation for an annual expense of amount X is calculated as $\min[\max(r \cdot X - d, 0), l]$. We also have the rating factors which reduces for pet insurance in France to specie, breed, age and gender. [Table 1](#) provides a list of the variables in the datasets.

Variable	Type	Description	Example
specie	character	Specie of the pet	dog
breed	character	Breed of the pet	australian sheperd
gender	character	Gender of the pet	female
insurance_carrier	character	identification number of the insurance company	1
age	numeric	Age of the pet (in years)	4 years
r	numeric	Value of the coverage rate	0.6
l	numeric	Value of the limit of the insurance coverage	1100
d	numeric	Value of the deductible of the insurance coverage	0
x	numeric	Yearly commercial premium	234.33

Table 1: List of the variables of our datasets

The first five rows are given in [Table 2](#).

specie	breed	gender	insurance_carrier	r	l	d	age	x
dog	australian sheperd	female	1	0.60	1100.00	0.00	2 years	221.34
dog	australian sheperd	female	1	0.70	1500.00	20.00	2 years	290.62
dog	australian sheperd	female	1	0.80	1800.00	30.00	2 years	361.53
dog	australian sheperd	female	1	1.00	2500.00	75.00	2 years	739.27
dog	australian sheperd	female	1	0.90	2200.00	50.00	2 years	594.28


Table 2: First five rows of our datasets

We have collected rates associated to 12 risk classes given in [Table 3](#).

Risk class #	specie	breed	gender	age
1	dog	australian sheperd	female	4 months
2	dog	australian sheperd	female	2 years
3	dog	australian sheperd	female	4 years
4	dog	french bulldog	female	4 months
5	dog	french bulldog	female	2 years
6	dog	french bulldog	female	4 years
7	dog	german sheperd	female	4 months
8	dog	german sheperd	female	2 years
9	dog	german sheperd	female	4 years
10	dog	golden-retriever	female	4 months
11	dog	golden-retriever	female	2 years
12	dog	golden-retriever	female	4 years

Table 3: The 12 risk classes under study

It means that we have 90 quotes to study each risk class. [Figure 9](#) provides a visual overview of the range of insurance coverage options available in the pet insurance market.



(a) Rates of coverage of the insurance policies (b) Deductible of the insurance policies (c) Limits of the insurance policies

[Figure 9](#): Overview of the insurance coverages offered by the five insurance companies operating in the French market under study.

We conduct two separate studies. In [Section 5.3](#), we focus on a specific risk class associated to a female, 4 years old, australian sheperd. Several claim models are compared. One is selected to look into the pricing strategies of the actors. In [Section 5.4](#), we look into the quotes of various risk classes that we investigate using a single model.

5.3 Analysis of one risk class using several models

We study 90 quotes from 5 insurers operating in the pet insurance market for a specific risk class associated to a 4-year-old female Australian Shepherd. We need to make some parametric assumptions to model claim frequency and severity. Classical claim frequency distributions include the Poisson, Binomial, and Negative Binomial distributions, which allow us to accommodate equidispersion, underdispersion, and overdispersion, respectively. For claim severity, we have chosen the Gamma and Lognormal distributions. The Gamma distribution is a common choice for modeling claim severity when using generalized linear models, but it is characterized by a light left tail. The Lognormal distribution has thicker tails, making larger claim sizes more likely to occur. We limit ourselves to two-parameter models: one parameter for claim frequency and another for claim severity. We consider three claim frequency distributions including $\text{Poisson}(\lambda)$, $\text{Bin}(12, p)$ and $\text{Geom}(p)$. The choice of setting the number of trials in the Binomial distribution to 12 aligns with our focus on annual expenses, making it a suitable choice to capture the monthly probability of a claim occurrence. The prior settings for the parameters are as follows:

$$\lambda \sim \text{Unif}([0, 10]), p \sim \text{Unif}([0, 1]). \quad (12)$$

We consider three claim severity distributions including $\text{LogNorm}(\mu = 0, \sigma)$, $\text{LogNorm}(\mu, \sigma = 1)$, and $\text{Gamma}(\alpha, \beta = 1)$. The prior settings over the parameters of the claim size distributions are as follows:

$$\mu \sim \text{Unif}([-10, 10]), \quad \sigma \sim \text{Unif}([0, 10]), \quad \text{and} \quad \alpha \sim \text{Unif}([0, 10^5]). \quad (13)$$

Combining the distributions for the claim frequency and severities results in a total of 9 loss models. The population size in the ABC algorithm is set to $J = 1,000$. The pure premiums are computed using $R = 2,000$ Monte Carlo replications. The algorithm stops whenever the difference between two consecutive tolerance levels is lower than $\Delta_\epsilon = 1$. The bounds for the loss ratio corridor are set to $\text{LR}_{\text{low}} = 40\%$ and $\text{LR}_{\text{high}} = 70\%$. The posterior distributions of the parameters for each model are provided in [Figure 10](#).




Figure 10: Posterior distribution of the parameters of the loss models when fitted to the pet insurance dataset.

For all the models, the algorithm updates the prior distribution in an informative way. [Table 4](#) provides the tolerance levels (ranked in increasing order) during the last iteration of the ABC algorithm for the loss models.

Model	ϵ
NegBin($b = 1, p$) – Gamma($\alpha, \beta = 1$)	87.50
Poisson(λ) – Gamma($\alpha, \beta = 1$)	93.72
Bin($a = 12, p$) – Gamma($\alpha, \beta = 1$)	94.69
NegBin($b = 1, p$) – LogNorm($\mu, \sigma = 1$)	95.85
Poisson(λ) – LogNorm($\mu, \sigma = 1$)	96.20
Bin($a = 12, p$) – LogNorm($\mu, \sigma = 1$)	96.24
NegBin($b = 1, p$) – LogNorm($\mu = 0, \sigma$)	108.64
Bin($a = 12, p$) – LogNorm($\mu = 0, \sigma$)	112.99
Poisson(λ) – LogNorm($\mu = 0, \sigma$)	113.15

Table 4: Tolerance level during the last iteration of the ABC algorithm fo each loss model

The final tolerance level for almost all the models lies between 87.5 and 113.15 which is higher than the tolerance obtained in the simulation study which was around 33 for 50 data points and 50 for 200 data points. This discrepancy indicates misspecifications which stem from our assumptions about insurance companies adhering to the expectation principle for premium calculation and the models employed for claim frequency and claim amounts. [Table 5](#) reports the estimations of the parameters of all the model using the MAP and the MODE.

[Table 6](#) reports the estimations of the average total claim amounts and the average loss ratio for all the models for all models when fitted using the MAP and the MODE.

We note that the risk level characterized here by the expected total claim amount is similar for all the models, maybe a bit higher for the models having the LogNorm($\mu = 0, \sigma$) as claim sizes distribution. We further look the loading function approximated via the isotonic regression. We estimate the pure premium for each model using the MAP as an estimator of the model parameters and we plot the isotonic regression function to explain the commercial premium on [Figure 11](#).

The isotonic fits of the loading function accross all the models are similar which means that the models all agree on a common ordering of the pure premiums of the various insurance coverages. To highlight the explanatory power of our methodology, let's focus on the Poisson(λ) – LogNorm($\mu, \sigma = 1$) loss model. Note that the choice of the loss model is somewhat arbitrary because the information extracted from the data in [Figure 11](#)

Model		MAP	MODE
Poisson(λ) – LogNorm($\mu, \sigma = 1$)	λ	0.31	0.30
	μ	6.14	6.19
Bin($a = 12, p$) – LogNorm($\mu, \sigma = 1$)	p	0.03	0.02
	μ	6.15	6.24
NegBin($b = 1, p$) – LogNorm($\mu, \sigma = 1$)	p	0.73	0.73
	μ	5.97	5.92
Poisson(λ) – Gamma($\alpha, \beta = 1$)	λ	0.16	0.14
	α	1365.53	1491.75
Bin($a = 12, p$) – Gamma($\alpha, \beta = 1$)	p	0.01	0.02
	α	1367.31	1191.38
NegBin($b = 1, p$) – Gamma($\alpha, \beta = 1$)	p	0.85	0.86
	α	1274.47	1260.02
Poisson(λ) – LogNorm($\mu = 0, \sigma$)	λ	5.30	4.24
	σ	2.97	3.11
Bin($a = 12, p$) – LogNorm($\mu = 0, \sigma$)	p	0.42	0.40
	σ	3.00	3.03
NegBin($b = 1, p$) – LogNorm($\mu = 0, \sigma$)	p	0.13	0.13
	σ	2.82	2.78

Table 5: MAP and MODE estimator for the parameters of the loss models.

Model	Loss ratio		E(X)	
	MAP	MODE	MAP	MODE
Poisson(λ) – LogNorm($\mu, \sigma = 1$)	0.62	0.62	239.34	245.04
Bin($a = 12, p$) – LogNorm($\mu, \sigma = 1$)	0.62	0.60	234.68	238.20
NegBin($b = 1, p$) – LogNorm($\mu, \sigma = 1$)	0.61	0.60	232.89	224.96
Poisson(λ) – Gamma($\alpha, \beta = 1$)	0.66	0.60	225.36	208.23
Bin($a = 12, p$) – Gamma($\alpha, \beta = 1$)	0.60	0.65	204.51	219.76
NegBin($b = 1, p$) – Gamma($\alpha, \beta = 1$)	0.63	0.64	220.44	221.00
Poisson(λ) – LogNorm($\mu = 0, \sigma$)	0.62	0.61	422.78	730.47
Bin($a = 12, p$) – LogNorm($\mu = 0, \sigma$)	0.64	0.64	510.01	564.42
NegBin($b = 1, p$) – LogNorm($\mu = 0, \sigma$)	0.60	0.60	322.39	299.85

Table 6: MAP and MODE estimators of the average loss ratio and average total claim amounts.

is relatively consistent across most of the considered models. In [Figure 12](#), we present a plot that illustrates the relationship between the commercial premium and the pure premium for the Poisson(λ) – LogNorm($\mu, \sigma = 1$) model. Different insurance companies are indicated by distinct colors, providing a visual representation of




Figure 11: Isotonic link between pure and commercial premium for the different loss models.

each company's respective rates.




Figure 12: Commercial premium as a function of the pure premium for the $\text{Poisson}(\lambda) - \text{LogNorm}(\mu, \sigma = 1)$ depending on the insurance carrier.

Fitting loss models enables us to condense the three-dimensional information of the rate of coverage, deductible, and limit into a single metric: the pure premium. Subsequently, isotonic regression unveils the relationship between commercial and pure premiums, providing a link between the two. The distinctions among various players in the pet insurance market come to light through the color-coded points, offering insights into the pricing strategies adopted by industry participants.

5.4 Analysis of several risk classes with one model

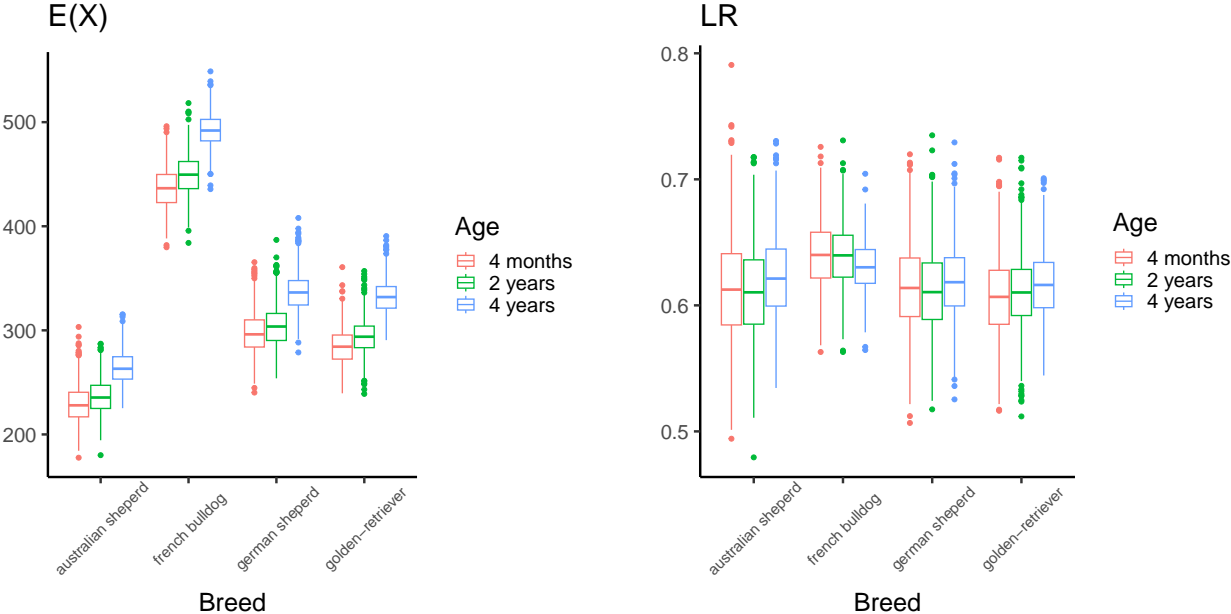
The $\text{Poisson}(\lambda) - \text{LogNorm}(\mu, \sigma = 1)$ model is fitted to the data within each risk classes (90 quotes) of Table 3. The prior settings are given by

$$\lambda \sim \text{Unif}([0, 10]), \text{ and } \mu \sim \text{Unif}([-10, 10]).$$

The algorithm’s hyperparameters are similar to that of the previous subsection with

$$J = 1000, R = 2000, \text{ and } \Delta_\epsilon = 1.$$

Figure 13 shows the posterior predictive distribution of the expected total claim amounts and the averaged loss ratio within each risk class.



(a) Posterior predictive distribution of $\mathbb{E}(X)$ (b) Posterior predictive distribution of the average loss ratio

Figure 13: Posterior predictive distribution of $\mathbb{E}(X)$ and average loss ratio within each risk class.

Figure 13a allows us to compare the different risk classes. We note that an older dog is more expensive on average and that the breeds may be ordered as Australian Sheperd, Golden-Retriever, German sheperd and french bulldog in terms of riskiness. Figure 13b indicates that the loss ratios are around 62 – 65% for all the risk classes.

6 Conclusion

We have developed a robust methodology for risk assessment based on market data. We employ a one-parameter model for the claim frequency and claim size distribution, connecting the pure premium to the commercial premium through an isotonic regression model. This approach optimizes the alignment between commercial and pure premiums while providing a framework for quantifying the associated parameter uncertainty through an Approximate Bayesian Computation algorithm.

The methodology's effectiveness and reliability have been validated within a simulation study and a practical application to an actual pet insurance dataset. This methodology is made accessible to the community through our R package, `IsoPriceR`².

While the results are promising, there remain avenues for further research. Future investigations can explore the selection of the most suitable model and consider the integration of historical data when it becomes available. One direction is the development of a credibility framework that combines historical and market data, providing a comprehensive perspective on risk assessment and pricing in emerging markets.

Acknowledgements

The authors express their gratitude to the four anonymous referees for their insightful comments, which significantly enhanced the quality of our manuscript. Pierre-O's work is conducted within the Research Chair DIALOG under the aegis of the Risk Foundation, an initiative by CNP Assurances.

References

- [1] Katrien Antonio and Emiliano A. Valdez. Statistical concepts of a priori and a posteriori risk classification in insurance. *AStA Advances in Statistical Analysis*, 96(2):187–224, February 2011. ISSN 1863-818X. doi: 10.1007/s10182-011-0152-7.
- [2] Richard E Barlow, HD Brunk, Daniel J Bartholomew, and James M Bremner. *Statistical inference under order restrictions. (the theory and application of isotonic regression)*. 1972.
- [3] Mark A Beaumont, Jean-Marie Cornuet, Jean-Michel Marin, and Christian P Robert. Adaptive approximate Bayesian computation. *Biometrika*, 96(4):983–990, 2009.

²see the [market_based_insurance_ratemaking](#) Github repository

- [4] Christopher Blier-Wong, H el ene Cossette, Luc Lamontagne, and Etienne Marceau. Machine learning in p&c insurance: A review for pricing and reserving. *Risks*, 9(1):4, dec 2020. doi: 10.3390/risks9010004.
- [5] Thomas A Dean, Sumeetpal S Singh, Ajay Jasra, and Gareth W Peters. Parameter estimation for hidden markov models with intractable likelihoods. *Scandinavian Journal of Statistics*, 41(4):970–987, 2014.
- [6] Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. An adaptive sequential monte carlo method for approximate bayesian computation. *Statistics and computing*, 22(5):1009–1020, 2012.
- [7] David Dickson. Principles of premium calculation. In *Insurance Risk and Ruin*, pages 38–51. Cambridge University Press, jan 2005. doi: 10.1017/cbo9780511624155.004.
- [8] Pierre-Olivier Goffard and Patrick J. Laub. Approximate bayesian computations to fit and compare insurance loss models. *Insurance: Mathematics and Economics*, 100:350–371, sep 2021. doi: 10.1016/j.insmatheco.2021.06.002.
- [9] Christian Gourieroux, Alain Monfort, and Eric Renault. Indirect inference. *Journal of Applied Econometrics*, 8(S1): S85–S118, dec 1993. doi: 10.1002/jae.3950080507.
- [10] Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029, July 1982. ISSN 0012-9682. doi: 10.2307/1912775.
- [11] Augustine Kong, Jun S. Liu, and Wing Hung Wong. Sequential imputations and Bayesian missing data problems. *Journal of the American Statistical Association*, 89(425):278–288, mar 1994. doi: 10.1080/01621459.1994.10476469.
- [12] Gareth Peters and Scott Sisson. Bayesian inference, monte carlo sampling and operational risk. *Peters GW and Sisson SA (2006)“Bayesian Inference, Monte Carlo Sampling and Operational Risk”*. *Journal of Operational Risk*, 1(3), 2006.
- [13] Gareth W Peters, Mario V W uthrich, and Pavel V Shevchenko. Chain ladder method: Bayesian bootstrap versus classical bootstrap. *Insurance: Mathematics and Economics*, 47(1):36–51, 2010.
- [14] Arthur E. Renshaw. Modelling the claims process in the presence of covariates. *ASTIN Bulletin*, 24(2):265–285, 1994. doi: 10.2143/AST.24.2.2005070.
- [15] Syoichi SASABUCHI, Manabu INUTSUKA, and D. D. Sarath KULATUNGA. A multivariate version of isotonic regression. *Biometrika*, 70(2):465–472, 1983. ISSN 1464-3510. doi: 10.1093/biomet/70.2.465.
- [16] Scott A Sisson, Yanan Fan, and Mark Beaumont. *Handbook of Approximate Bayesian Computation*. Chapman and Hall/CRC, 2018.
- [17] Mario V. W uthrich and Johanna Ziegel. Isotonic recalibration under a low signal-to-noise ratio, 2023.

A Comparison of linear and isotonic regression to predict commercial premiums

Instead of the linear link of (14) between pure and commercial premium, we consider

$$\pi_i = a_i \exp(-b_i \exp(cp_i)), \text{ for } i = 1, \dots, n. \tag{14}$$

where

$$a_i \sim \text{Unif}(5, 10), \quad b_i \sim \text{Unif}(2, 6), \quad \text{and } c = 2 \text{ for } i = 1, \dots, n.$$

This is a Gompertz growth curve type of link. The commercial premiums as a function of the pure premium is shown on Figure 14.




Figure 14: Isotonic link between the pure and commercial premiums.

We further compare the residuals of the isotonic regression model fitted to the data of Figures 3 and 14 to that of a linear regression model fitted to the same data on Figure 15. We note the proximity of the two models when the link between the pure and commercial premium is linear. When the link is not linear then isotonic regression model outperforms linear regression.

B Other premium principle

This paper focuses on the expectation premium principle as we try to inform the link f between the commercial premium π and the pure premium $p = E[g(X)]$. Other premium principles such as the standard deviation




Figure 15: Boxplot of the residuals of the linear and isotonic regression models fo a linear and a Gompertz type link between pure and commercial premiums.

principle can be considered by slightly adapting the method. Under such principle we have

$$\pi = f\left(\mathbb{E}[g(X)], \sqrt{\mathbb{V}[g(X)]}\right). \quad (15)$$

where $f : \mathbb{R}_+ \times \mathbb{R}_+ \mapsto \mathbb{R}_+$. The same methodology applies, we simply need a bivariate modle for f . The commercial premium should be increasing whenever the pure premium or the variance of the risk increases which leads to consider generalization of the univariate isotonic regression models which are readily available in the litterature see the work of SASABUCHI et al. [15]. More sophisticated premium principles such as the Escher principle or the utility indifference principle are also possible. Premium principles are described at length in actuarial science textbooks such as Dickson [7]. Considering a premium principle instead of another leads to model misspecification and will impact the final estimates of the underlying risk.