Quantile Mixing and Model Uncertainty Measures

Thierry Cohignac, Nabil Kazi-Tani

Quantile Mixing and Model Uncertainty Measures

Thierry Cohignac* Nabil Kazi-Tani†

December 11, 2019

Abstract

In this paper, we introduce a new simple methodology for combining two models, which are given in the form of two probability distributions. We use convex combinations of quantile functions, with weights depending on the quantile level. We choose the weights by comparing, for each quantile level, a given measure of model uncertainty calculated for the two probability distributions that we want to combine. This methodology is particularly useful in insurance and reinsurance of natural disasters, for which there are various physical models available, along with historical data. We apply our procedure to a real portfolio of insurance losses, and show that the model uncertainty measures have a similar behavior on the set of various insurance losses that we consider. This article serves also as an introduction to the use of model uncertainty measures in actuarial practice.

Key words: Model combination, Model uncertainty, Quantiles, Risk management, Catastrophe models.

AMS 2010 subject classifications: 91B30, 91G70.
JEL codes: C52, C44, C53, C15.

Contents

*Thierry Cohignac, CCR, 157, Boulevard Haussmann, 75008 Paris - France, *Email:*tcohignac@ccr.fr
†Nabil Kazi-Tani, ISFA, Université de Lyon 1, 50 Avenue Tony Garnier, 69007 Lyon, France *Email:* nabil.kazi-tani@univ-lyon1.fr.

1 Introduction

There are various practical situations in which one is confronted to several models for the same phenomenon. A natural problem is then to choose among these models, or to combine them to take decisions. This is true in particular in the domain of insurance and reinsurance of natural catastrophes, where one has access to several Catastrophe models (Cat models) available in the market, but also potentially to models which are developed internally by a given insurance or reinsurance company. These models rely on the physical modelling of natural phenomenons, such as flooding, earthquakes or wind speeds for example. These tools are also in competition with models that originate only from historical data, but which can be rather scarce for natural catastrophes. Combining these sources of information is an important problem, in particular since natural catastrophes represent a major issue for insurers and reinsurers, causing losses up to 140 billion dollars in 2017 (see [20]).

There is a large literature on models combination: one of the main methodology is the *linear pool*, which consists in taking a linear combination of each model's results. It is a standard assumption that the output of each model is given in the form of a probability measure. Most of the literature on linear pooling deals with convex combinations of the densities provided by the models or of the cumulative distribution functions (CDF). In the present paper, we combine probability measures by taking convex combinations of the associated quantile functions. Indeed, quantile functions, just as densities of CDF, characterize the distribution of a random variable. In the present paper, we provide a procedure to choose the weights, which is based on model uncertainty measurements.

Taking convex combinations of quantile functions allows to make the weights of the convex combination depend on the quantile level, which is easily interpretable in practice. Indeed, we have at our disposal both historical data and a Cat model, and the criteria that we introduce in this paper to choose the weights, which is based on model uncertainty measures, leads to hight weights for the historical model quantile for frequent claims and similarly hight weights for the Cat model quantile for rare and extreme claims. One drawback of this approach is that the curve that we build is not necessarily a quantile function, for instance it can happen that it is not increasing: note however that the recent results from [4] allow to deal exactly with this type of situation, using techniques from optimal transport and rearrangement.

Summary of the approach

Assume that the models that we need to combine are given in the form of two probability measures ν_h and ν_e on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$, where $\mathcal{B}(\mathbb{R})$ is the Borel σ-algebra on $\mathbb{R}$. Let Q_h and Q_e be the quantile functions respectively associated to the measures ν_h and ν_e (see (2.1) for a precise definition), and consider the function Q given by

$$Q(u) := \lambda(u)Q_e(u) + [1 - \lambda(u)]Q_h(u), \tag{1.1}$$

where $\lambda : [0, 1] \to [0, 1]$ is a weight function. To the best of our knowledge, using a model combination in the form of (1.1) was first suggested by [13] (equation (7)). Additional information available on ν_h and ν_e may lead to a priori information on the curve $u \mapsto \lambda(u)$: for instance, if we know that ν_h is constructed from historical data, while ν_e comes from a physical Cat model, which is a model based on the underlying physical natural hasard (for example floods), then it is natural to expect that $\lambda(u)$ will be hight for small values of u and small for large values of u. Indeed, small values of u correspond to more frequent claims, for which it is expected to have larger samples of historical data.

To compute the weight function λ, we compare model uncertainty measures associated to ν_h and ν_e, which have been introduced in [2]. Given a set of insurance portfolios and associated data set, the methodology that we introduce for choosing the weight $\lambda(u)$ appearing in front of $Q_e(u)$ in (1.1), consists in computing the proportion of insurance portfolios for which the model uncertainty associated to ν_e is lower than the model uncertainty associated to ν_h (see (2.5)). This simple approach is easily interpretable.

Related literature

As already mentioned above, there is a large literature on models combination, and several survey papers are available, such as [10], [13], or [5], we also refer the reader to Chapter 14 of the recent book [8]. The general idea behind model combination is that when faced with multiple models for the same phenomenon or for the same outcome, averaging is likely to improve accuracy, compared to choosing a single model. If a model is represented as a probability measure, which is the point of view that we adopt in this paper, models can be combined using their related density functions, cumulative distribution functions (CDF), or any functional characterising the distribution. Densities or CDFs combinations seems to be the most standard approach in the literature, and the most common way to combine them is the *linear pooling* which consists in taking linear or convex combinations. We refer the reader to [17] for a general result on the linear pooling.

Non-linear combination techniques have also been studied, and usually the goal is to maximise a given score function (see [11], [12] and the references therein). In this paper, we still exploit the properties of the simple linear combination methodology, but we apply it to quantile functions. Notice that this methodology, seen as a function of CDFs, which are just the generalized inverses of quantiles, is a highly non linear operation.

The other important question related to this paper is *model uncertainty*, which has recently seen a tremendous rise of popularity among academics and professionals. This question is not new however, and it can be considered at the heart of robust statistics [15] or stochastic control theory [9]. Recent papers ([7], [19], or [14]) have focused on volatility uncertainty in utility maximisation models, in pricing and hedging of financial derivatives or in the study of stochastic differential equations.
In our specific case, we are working in a time static framework, and since we are interested in quantiles, it is natural to look at model uncertainty related to risk measurement. This paper relies on results of [2], where the authors introduce *numerical* measures of model risk, associated to a given risk measure. The aim of the present paper is to use these measures of model risk to combine quantiles.

Main contributions

The current paper introduces a new way to combine two probability distributions, through the convex combination of their associated quantiles, with weights computed by comparing model uncertainty measures. The issue of combining probability distributions is important in insurance and reinsurance of natural disasters.
As a byproduct, this paper introduces the use of model uncertainty measures in actuarial practice.
Lastly, our tests on a real portfolio show that the difference between model uncertainty measures corresponding to losses computed on historical data and Cat softwares behave similarly for a relatively large class of insurance data. We think that this empirical result is interesting on its own and could be tested on other real data portfolios.

Structure of the paper

In Section 2, after briefly recalling some well known needed definitions and properties of VaR and TVaR risk measures, we introduce the model uncertainty measure that we use, and how we apply it to construct a weighting curve for our quantile combination. Then in Section 3, we show how the computed model uncertainty measures compare, what type of weighting curve it gives in practice, and the form of combined quantiles that we get.

2 Model and Assumptions

In this Section, we introduce our methodology for choosing the weight function λ in (1.1). Let us first recall the formal definitions and some properties of the Value-at-Risk (VaR) and Tail-Value-at-Risk that will be needed later.

2.1 VaR and Tail-VaR

Let $(\Omega, \mathbb{F}, \mathbb{P})$ be a fixed probability space. For a given random variable X, F_X denotes its CDF defined by $F_X(x) := \mathbb{P}(X \leq x)$ and Q_X denotes its quantile function, defined as the generalized inverse of F_X:

$$Q_X(u) := \inf\{x \in \mathbb{R} \mid F_X(x) \geq u\}. \tag{2.1}$$

Notice that Q_X only depends on the distribution of X. If X has distribution μ, we will also write Q_μ for the quantile function associated to X. The well known Value-at-Risk is then just defined as a quantile:

Definition 2.1. *The Value-at-Risk (VaR) at level $\alpha \in (0,1)$ associated to X is given by $VaR_\alpha(X) := Q_X(\alpha)$.*

This measure is widely spread in insurance and reinsurance practice, as it is easy to interpret and to implement than most other measures of risk. Another reason lies in its links with European solvency regulations: indeed, the Solvency II rules retain the VaR as a risk measure for regulatory capital calculations. The main criticism around the VaR concentrates around its lack of subadditivity. Nevertheless, there are large classes of distributions within which the VaR is subadditive, we refer the interested reader to [6] for more details and a discussion.

For a random variable X with a continuous CDF, we have $\mathbb{P}(X \leq VaR_\alpha(X)) = \alpha$. However, the VaR does not provide any information on the event $\{X > VaR_\alpha(X)\}$. Another well known measure, which takes this aspect into account is the Tail Value-at-Risk, defined as follows.

Definition 2.2. *The Tail Value-at-Risk (TVaR) at level $\alpha \in (0,1)$ associated to X is given by*

$$TVaR_\alpha(X) := \frac{1}{1-\alpha} \int_\alpha^1 VaR_u(X) du.$$

In the case where X has a continuous CDF, TVaR coincides with the Conditional Tail Expectation (CTE), which is the average loss conditioned on the event $\{X > VaR_\alpha(X)\}$:

$$CTE_\alpha(X) := \mathbb{E}[X \mid X > VaR_\alpha(X)].$$

TVaR is also a popular risk measure in reinsurance practice, and satisfies a set of properties, including subadditivity, that makes it a coherent risk measure in the sense of [1].

In this paper, we use extremal properties of VaR and TVaR to construct a weighting curve in a quantile mixing framework.

2.2 Model Uncertainty Indexes

The actual calculation of the risk measures described above requires the prior estimation of the distribution of the underlying risk X. As a result, the choice of the risk distribution is subject to an uncertainty which weighs on the calculation of the risk measure.
In the present paper, we rely on the results of [2], that provides several different quantitative measures of model risk. One of the ideas of the authors in [2] is to consider a "worst case" approach in which oneself is placed in the most unfavorable case: given an acceptable set of probability distributions, we are interested in the maximal value over the acceptable distributions (or the minimum value, if we consider gains rather than losses). The knowledge of the first moments of the risk distribution provides one typical example of acceptable set.

Let ρ be a given risk measure. We will use the following quantitative model risk measurement proposed in [2], known as the *absolute measure of model risk*, and defined by

$$AM(X_0, \mathcal{L}) := \frac{\overline{\rho}(\mathcal{L}) - \rho(X_0)}{\overline{\rho}(\mathcal{L})} = 1 - \frac{\rho(X_0)}{\overline{\rho}(\mathcal{L})}, \tag{2.2}$$

where $\mathcal{L}$ is a given set of probability distributions, X_0 is a random variable with a reference distribution in $\mathcal{L}$, and where

$$\overline{\rho}(\mathcal{L}) := \sup_{X \in \mathcal{L}} \rho(X).$$

Notice that the definition that we give above is different from the original definition given in [2], where the absolute measure of model risk AM^* defined there is such that $AM^* = \frac{AM}{1-AM}$. Both definitions make use of a reference distribution X_0. The distribution of X_0 being given, $AM(X_0, \mathcal{L})$ measures the rate of decrease to go from the supremum $\overline{\rho}(\mathcal{L})$ to the particular value $\rho(X_0)$. In particular $AM(X_0, \mathcal{L})$ is real number lying between 0 and 1. In comparison, AM^* measures the rate of increase to go from $\rho(X_0)$ to $\overline{\rho}(\mathcal{L})$, which is non negative but can be greater than 1.

$AM(X_0, \mathcal{L})$ is a measure of model uncertainty in the sense that $AM(X_0, \mathcal{L}) = 0$ implies that $\rho(X_0) = \overline{\rho}(\mathcal{L})$. In that case, the capital to put in reserve cannot exceed the value $\rho(X_0)$, and we interpret this as no model uncertainty. When $AM(X_0, \mathcal{L}) = 1$, then $\rho(X_0) = 0$ and we interpret this as maximal model uncertainty. Notice that to obtain a model uncertainty measure between 0 and 1, one can also use the *relative measure of model uncertainty*, introduced in [2] and defined by

$$RM(X_0, \mathcal{L}) := \frac{\overline{\rho}(\mathcal{L}) - \rho(X_0)}{\overline{\rho}(\mathcal{L}) - \underline{\rho}(\mathcal{L})}$$

where

$$\overline{\rho}(\mathcal{L}) := \sup_{X \in \mathcal{L}} \rho(X) \quad \text{and} \quad \underline{\rho}(\mathcal{L}) := \inf_{X \in \mathcal{L}} \rho(X).$$

The definition of $RM(X_0, \mathcal{L})$ makes use of both the supremum and infimum values of ρ. The definition (2.2) that we use here is asymmetric and concentrates on extremal high values, i.e. on the supremum.

Let $\mu \in \mathbb{R}$, $\sigma > 0$, $-\infty \leq A$, $B \leq +\infty$ and let $\mathcal{L}(A, B, \mu, \sigma)$ be the set of probability measures with support $[A, B]$ and mean and variance respectively given by μ and σ^2. Notice that A and B can take infinite values. The interesting case for us in the reinsurance applications will be $A = 0$ and $B = +\infty$. The paper [16] provides the value of $\overline{\rho}(\mathcal{L}(A, B, \mu, \sigma))$, when ρ is either the Value-at-Risk or the Tail Value-at-Risk.

Proposition 2.1. *Let μ and σ be fixed and let $\mathcal{L}^+ := \mathcal{L}(0, +\infty, \mu, \sigma)$. Then the following equality is correct for both $\rho = VaR_\alpha$ and $\rho = TVaR_\alpha$*

$$AM(X_0, \mathcal{L}^+)(\alpha) = \begin{cases} 1 - \frac{\rho(X_0)}{\mu + \sigma\sqrt{\frac{\alpha}{1-\alpha}}} & \text{if} \quad \alpha^* \leq \alpha \leq 1 \\ 1 - (1-\alpha)\frac{\rho(X_0)}{\mu} & \text{if} \quad 0 \leq \alpha \leq \alpha^*. \end{cases} \tag{2.3}$$

where $\alpha^* := \frac{\sigma^2}{\mu^2+\sigma^2}$.

Proof. This is an immediate application of Theorem 3.1 in [16]. □

Remark 2.1. *Observe that the value of $AM(X_0, \mathcal{L}^+)$ given in (2.3) is a continuous function of α. Indeed for $\alpha = \alpha^*$ both values in (2.3) are equal to $1 - \mu\frac{\rho(X_0)}{\mu^2+\sigma^2}$. Notice also that $\rho(X_0)$ is a function of α, wether ρ is the VaR_α or $TVaR_\alpha$ risk measure. In particular, the quantity AM above is not a linear function of α when $\alpha \leq \alpha^*$. When no confusion is possible, we will write $AM(X_0, \mathcal{L}^+)$ instead of $AM(X_0, \mathcal{L}^+)(\alpha)$ for simplicity of notations.*

Remark 2.2. *When the set $\mathcal{L}$ contains only random variables with range equal to $\mathbb{R}$, i.e. when $A = -\infty$ and $B = +\infty$, then $AM(X_0, \mathcal{L}^+)$ is equal to the first value in (2.3) for every $\alpha \in [0, 1]$ (i.e. $\alpha^* = 0$).*

Let us now describe how the values $AM(X_0, \mathcal{L}^+)$ can be used to combine different possible models for the same underlying phenomenon.

2.3 Constructing a weighting curve

For a given portfolio of risks related to natural disasters, we can estimate the losses distribution using either available historical data, or using a Cat model that simulates random natural events. The idea of the methodology that we introduce is to compute a model risk measure for both the reference model coming from historical data and the reference model coming from a Cat software. As one would expect, it turns out that the historical model performs better for quantile levels close to 0 and by opposition, the Cat software performs better for large quantile levels, that is to say for extreme risks.

More precisely, let ν_h and ν_e be the distributions estimated respectively from historical data and from a Cat software (the subscript e in ν_e stands for *exposition*, which is a standard terminology to designate Cat models based on the physical modelling of natural hazards). Let $\alpha \mapsto AM_h(\alpha)$ and $\alpha \mapsto AM_e(\alpha)$ be the absolute measure of model risk given in (2.3) when X_0 has a distribution respectively given by ν_h and ν_e.
Figure 1 below displays the values of AM_h (red curve) and AM_e (green curve), for α varying between 0 and 1, constructed using real natural disaster losses data from a given insurance company up to 2017. The model ν_e corresponds to the output of the internal Cat model developed by CCR, run on a given insurance company portfolio. See Section 3 for more details on how we computed the model risk values. For obvious confidentiality reasons, we will not be more specific about the data we use, but for the sake of research reproducibility, we indicate in detail the procedure we used to compute the model uncertainty measures, so that the empirical results in this paper can be tested on other real data sets. Notice that these two curves cross only once, and that $AM_e(\alpha) \leq AM_h(\alpha)$ for α greater than a value approximately equal to 0.98. Figure 2 is a zoom on large values of α.
From Remark 2.1, we know that these curves are continuous, and the breakpoint α^* is clearly visible on Figure 1. It is equal approximately equal to 0.22 for the historical model ν_h and to 0.65 in the

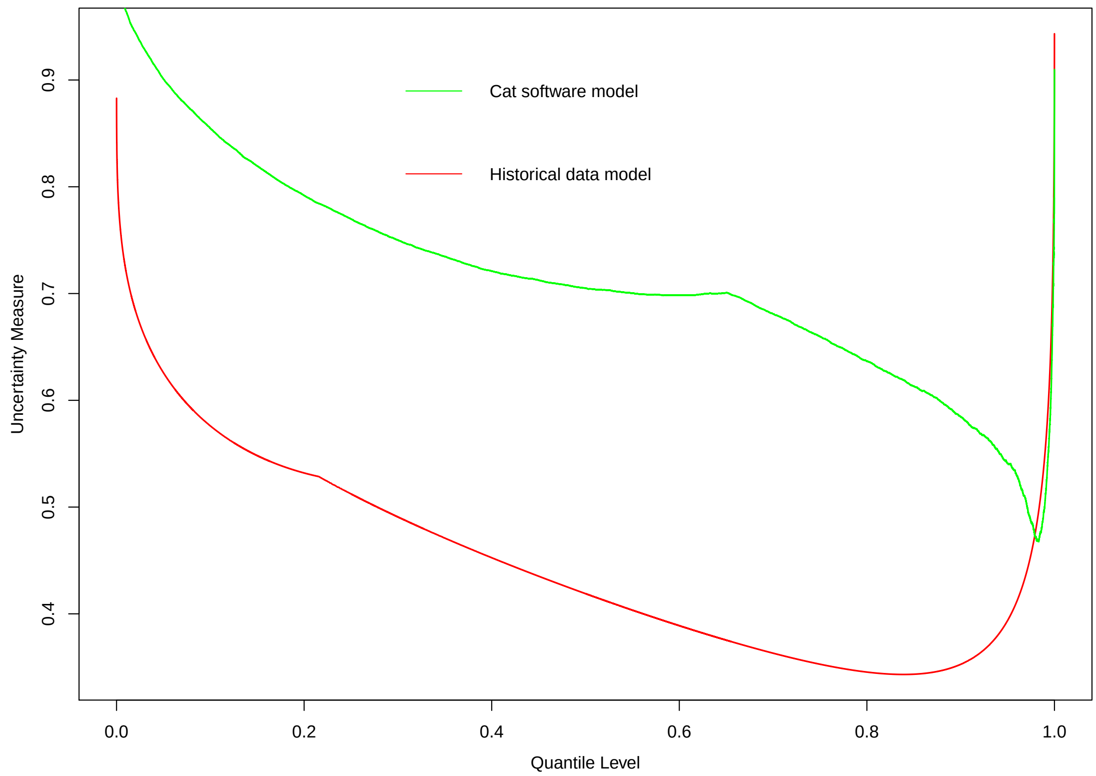

Figure 1: Illustration of the measure of model risk defined in 2.2.

case of ν_e. One can already see from here that even if ν_h and ν_e are two estimated models for the same phenomenon, their first two moments can be rather different.
Assume that we have at our disposal the historical data and portfolio information of n insurance companies, so that we are able for each of these companies to construct the associated measures ν_h and ν_e. Let $\nu_h^{(i)}$ and $\nu_e^{(i)}$ be respectively the estimated historical and Cat software model for insurer i, and let $AM_h^{(i)}$ and $AM_e^{(i)}$ be the associated model uncertainty measures. Assume also that for each i between 1 and n, there exists a value $\alpha_i \in (0,1)$ such that $AM_h^{(i)}(\alpha) \leq AM_e^{(i)}(\alpha)$ for $\alpha \leq \alpha_i$ and $AM_h^{(i)}(\alpha) \geq AM_e^{(i)}(\alpha)$ for $\alpha \geq \alpha_i$. This is the case for instance for the data in Figure 1, and it will also be the case in our illustrations in Section 3.

Let us now describe our methodology to construct a weighting curve λ, as it appears in (1.1). Let Q_h and Q_e be the quantile functions respectively associated to ν_h and ν_e and suppose that we want to combine these two models using a convex combination of their quantile functions as follows:

$$Q(u) := \lambda(u)Q_e(u) + [1-\lambda(u)]Q_h(u), . \tag{2.4}$$

The main idea is simple: since $\lambda(u)$ is the weight associated to the Cat software model, we want λ to be high where the Cat software model performs better, that is to say where its associated measure of model risk is lower. Thus we define

$$\lambda(u) := \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}_{\{\alpha_i \leq u\}}. \tag{2.5}$$

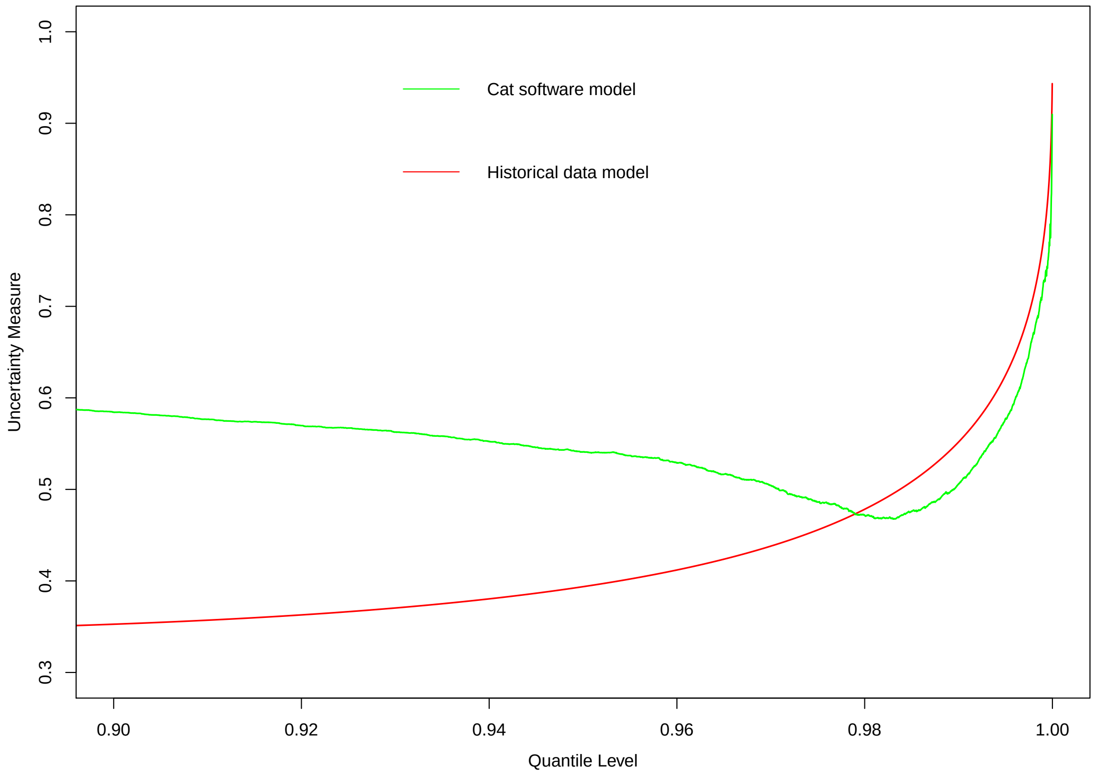

Figure 2: Zoom on α values between 0.9 and 1.

For a fixed u, $\lambda(u)$ gives the proportion of insurance companies for which the Cat software model performs better in terms of model uncertainty. Notice that this can be rewritten:

$$\lambda(u) := \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}_{\{AM_e^{(i)}(u) \leq AM_h^{(i)}(u)\}}.$$

Remark 2.3. *The map $u \mapsto \lambda(u)$ defined in* (2.5) *is also the cumulative distribution function of the probability measure on the finite set $\{\alpha_1, \ldots, \alpha_n\}$ with uniform weights. This interpretation allows possible generalizations of* (2.5)*, for example by using non uniform weights for the α_i's. This can be done using some significance criteria for each insurance company i, for example its relative size in the portfolio, or the size of the available data.*

3 Numerical illustration

3.1 Model Uncertainty Indexes for the Reference Models

As shown in (2.5), the weighting curve that we construct is based on the comparison between the model uncertainty indexes computed from both the historical and the Cat software models.

To construct the historical model, we apply standard parametric inference procedures on the available data sets. We end up with the estimated distribution ν_h. To compute AM_h, we need to estimate $\rho(X_0)$, when X_0 has distribution ν_h. When ρ is the Value-at-Risk, this is done either by explicitly inverting the CDF (in the case of the Pareto or Lognormal distributions for instance) or by numerically inverting the CDF.

On the other hand, the measure ν_e is obtained as the output of a Cat software: in the Appendix below, we provide a quick summary on the way the software generates simulation years. For the illustrations below, we used a set of 50000 simulated years for each given insurance portfolio i. We took the corresponding empirical measure as our estimate for ν_e. In particular, to compute AM_e, we need the value of $\rho(X_0)$ when X_0 has distribution ν_e, and when ρ is the Value-at-Risk, we estimated it using the standard order statistics estimator.

3.2 The Weighting Curve

Suppose that ρ is the Value-at-Risk at level α. We analyzed a set of 93 real insurance portfolios, among which $n = 77$ (83%) have a model uncertainty profile as described in Figures 1 and 2, i.e. the corresponding curves AM_h and AM_e only cross once, with the Cat software model performing better than the historical model for large values of the quantile level α. We kept thus a set $\alpha_1, \ldots, \alpha_{77}$ of crossing points. Figure 3 below shows the obtained curve $u \mapsto \lambda(u)$ as defined in 2.5.

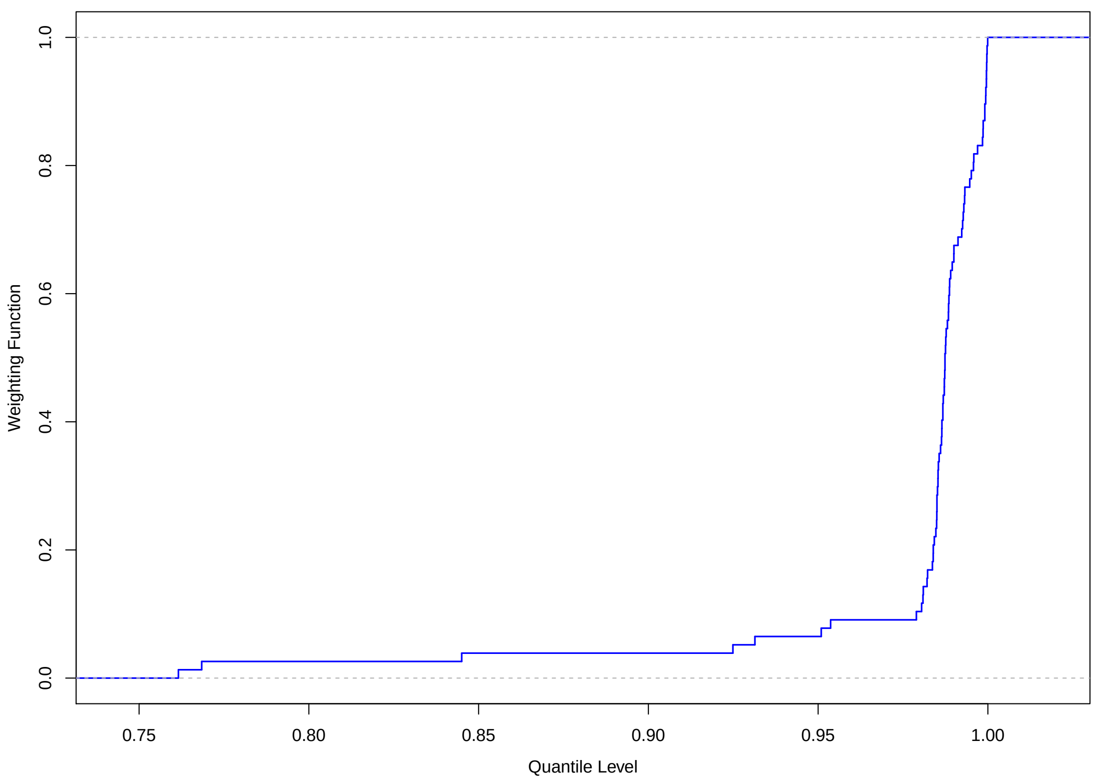

Figure 3: Zoom on α values between 0.9 and 1.

One can see that the values are rather concentrated around $u = 0.98$, corresponding to a return period equal to $R = \frac{1}{1-u} = 50$ years. Indeed the empirical average of the vector $\alpha_1, \ldots, \alpha_{77}$ is equal to 0.9794 and the empirical standard deviation is 0.0408.

To add robustness to this weighting curve, it is possible to make it more regular. Indeed, a more regular curve avoids abrupt changes from one year of data to the other. To proceed with this regularization, we chose to apply a given kernel K to the weighting curve. More precisely, recall from Remark 2.3 that the map $u \mapsto \lambda(u)$ can be interpreted as the CDF of a probability measure. It is possible to estimate this CDF using a standard kernel estimator as defined in [18]. Let K be a given kernel function, i.e. a non-negative real valued integrable symmetric function, whose integral over $\mathbb{R}$ is normalized to 1. Then, using a CDF kernel estimator, formula (2.5) for the weighting curve becomes:

$$\lambda_K(u) := \frac{1}{n}\sum_{i=1}^{n} \mathcal{K}_H(u - \alpha_i), \tag{3.1}$$

where $\mathcal{K}_H(u) := \mathcal{K}(H^{-1/2}u)$ is the scaled integrated kernel, with $\mathcal{K}(u) := \int_{-\infty}^{u} K(w)dw$, and H represents the bandwidth size.

In figure 4, we show the obtained regularized weighting curve, with the choice $H = 0.02$ and where K is the standard Gaussian density. $H = 0.02$ corresponds to half the standard deviation value. We tested the different curves obtained with H between 0.01 and 0.05 and this does not have a significant impact on the use we make of the obtained quantiles (for pricing purposes for instance). This means, at least for the particular application we make with the quantiles we obtain, that the overall shape of the weighting curve matters, rather than the particular choice of the bandwidth H.

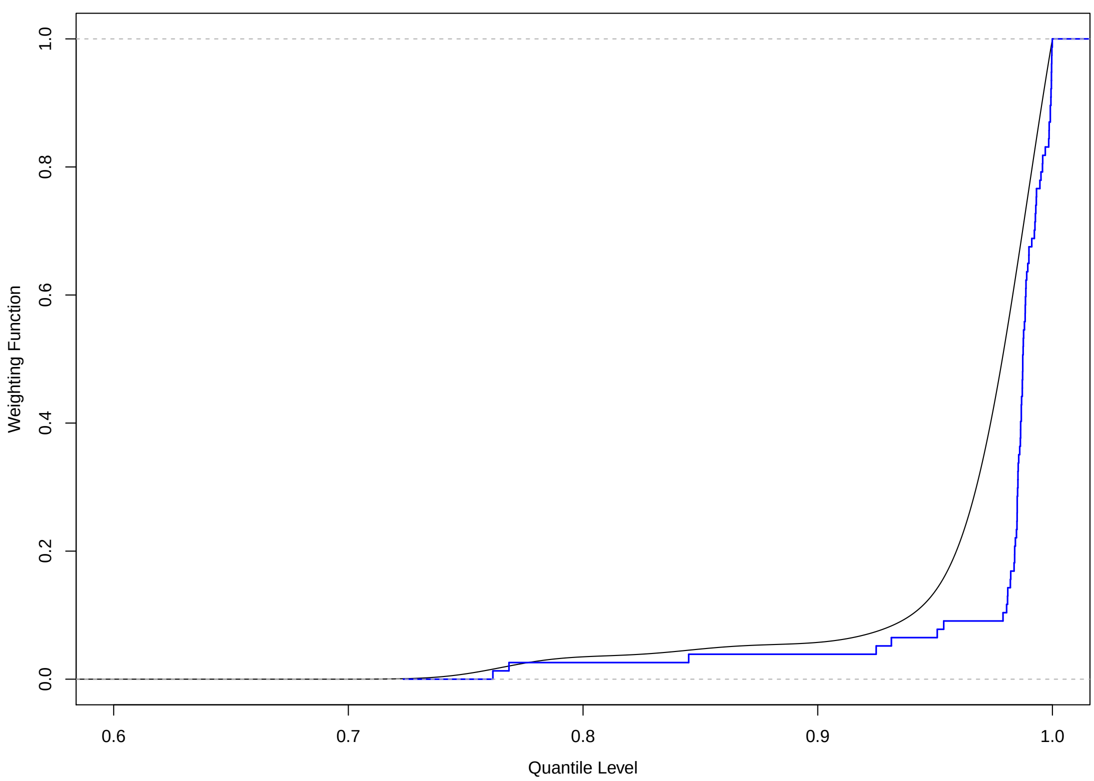

Figure 4: Regularized weighting curve λ_K defined in (3.1).

3.3 Obtained quantiles

Now that we have at our disposal a mixing curve λ, or its regularized version λ_K, we can use it to combine probability measures via equation (2.4). Figure 5 below shows the quantile functions that we get for the data associated to a given insurance company, using the Cat software model, the historical data, and the combined model. We only plotted the values for quantile levels $u \geq 0.8$, since for $u \leq 0.9$, the combined model is almost equal to the one constructed from historical data. Said otherwise, for events with a return period lower than 10 years, the historical model only is used. Then, the weight associated to the Cat model increases until the combined model is almost equal to the Cat software model.

Recall that we have at our disposal losses data from a portfolio of different insurance companies. Recall also that the definition (2.5) of the mixing curve λ makes use of the data from a large majority of these insurance companies. Thus the curve λ is constructed at a macro level, but we use this mixing curve to combine the quantile functions at a micro level, i.e. for individual insurance companies, as well as at a macro level, i.e. for the whole portfolio quantile.

Remark 3.1. *There is no reason guaranteeing that the function Q defined in* (2.4) *in non-decreasing. So* (2.4) *does not define a quantile function, even if in all our numerical computations, the function Q that we obtain is increasing. If the monotonicity property of the quantile is violated, one can use the methodology described in [4], based on optimal transport arguments, consisting in rearranging the obtained function Q into a non-decreasing function.*

3.4 The case of the Tail Value-at-Risk

Let us now give more details on the procedure we employed when $\rho = TVaR_\alpha$, for $\alpha \in (0,1)$.

First of all, notice that to compute $TVaR_\alpha$, we need to numerically evaluate the integral appearing in Definition 2.2. To this end we used Simpson's rule, which consists in approximating the value of the integrand by a quadratic polynomial on each point of the interval subdivision. This method is known to have good numerical precision (see for instance [3]). More precisely, we used the following approximation:

$$TVaR_\alpha(X) \approx \frac{1}{6n}\left(Q_X(\alpha) + Q_X(1) + 2\sum_{i=1}^{n-1} Q_X(u_i) + 4\sum_{i=0}^{n-1} Q_X(\frac{u_i + u_{i+1}}{2})\right),$$

where $(u_0, \ldots, u_n)$ is a uniform subdivision of the interval $[\alpha, 1]$ with step $\frac{1}{n} = 10^{-5}$. Now to compute the model uncertainty indexes, we need to compute $TVaR_\alpha(X_0)$, when the distribution of X_0 is computed both from historical data and from the Cat software model. To do so, in the $TVaR_\alpha$ approximation above, we used the same quantile estimators as those described in subsection 3.2.

Now that we are able to compute $TVaR_\alpha$ values, we can use this to evaluate the model uncertainty indexes when the reference risk measure is the Tail VaR, and then construct an associated smoothed weighting curve λ_K, and use the weighting curve to produce quantile values, as we did in the previous subsection.
In figure 6 below, we plotted the obtained weighting curves λ and λ_K for $H = 0.02$. Compared to Figure 4, the crossing points between model uncertainty measures are more dispersed. As a result, this curve gives more weight to the Cat software model.

Remark 3.2. *Obviously, when $TVaR_\alpha(X)$ is computed, all the quantile values $Q_X(u)$ are used, for u between α and 1. So when we compute the model uncertainty measure given in* (2.3) *with $\rho = TVaR_\alpha$,*

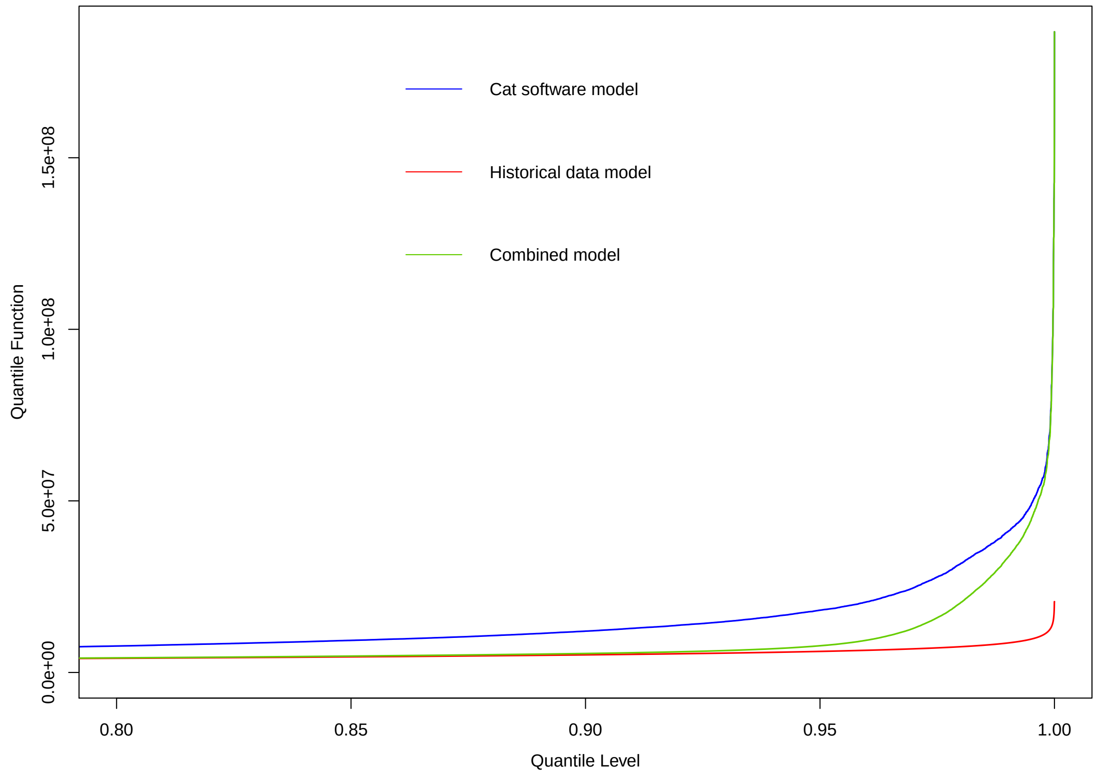

Figure 5: Quantile functions obtained with the Cat software model, the historical data, and the combined model defined in (2.4).

with α *small, we are using almost all the distribution function associated to* X*. This is in contrast to the calculation of the quantity in* (2.3) *when* $\rho = VaR_\alpha$*, which mainly explains he differences that we obtain for the weighting curves computed with* $TVaR_\alpha$ *or* VaR_α *as the reference risk measure.*

Figure 7 below shows the obtained quantile values, for the same real data set used to produce Figure 5. The blue and red curves on these two graphs are the same, one can observe however that the green curve corresponding to the combined model, has slightly higher values, since it gives more weight to the Cat software model.

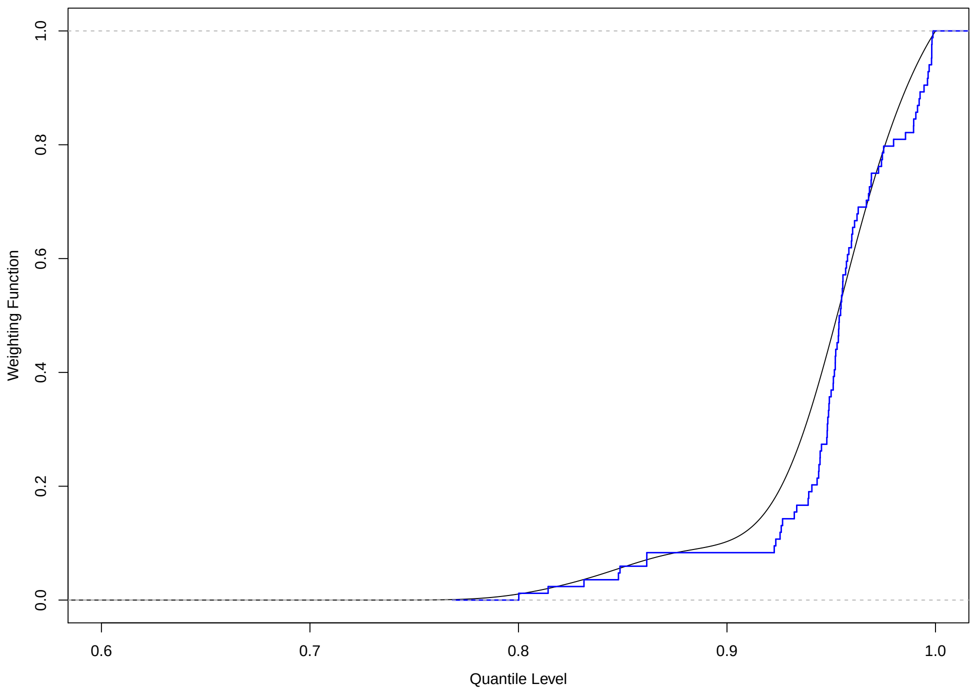

Figure 6: Smoothed weighting curve λ_K defined in (3.1), with $H = 0.02$ and $\rho = TVaR$.

Appendix

We made reference in this paper to Cat software models. The main output of these softwares is given in the form of an Event Loss Table (ELT), which can be used to produce losses simulations. Since AM_e is calculated using these simulations, let us briefly explain in this Appendix how they are constructed. We start with the following simplified ELT:

Events	Frequency	Claim distribution
1	γ_1	μ_1
2	γ_2	μ_2
$\dots$	$\dots$	$\dots$
M	γ_M	μ_M

The first column contains the numbers of M natural events whose physical parameters are fixed. These events are considered by the software as the exhaustive list of events that can occur in the geographical zone within which lie the risks of the input portfolio (M is typically of the order of several thousands).
The third column can be read as follows: if natural event number i occurs, then the cost for the input insurance portfolio is distributed according to μ_i.

We can then draw simulations in the following way: we first simulate a Poisson variable N with parameter γ, where $\gamma := \gamma_1 + \cdots + \gamma_M$. Once we have a value $N(\omega)$, we draw $N(\omega)$ independent

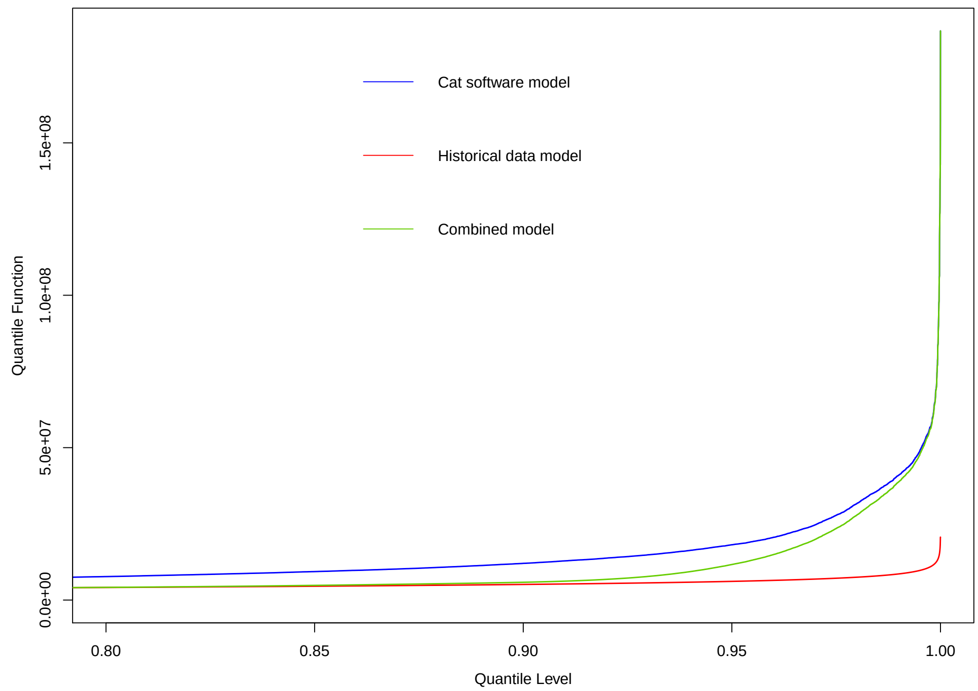

Figure 7: Quantile functions obtained with the combined model defined in (2.4) and with $TVaR$ as the reference risk measure.

integers in $\{1, \ldots, M\}$ with distribution $\{p_1, \ldots, p_M\}$ with

$$p_i := \frac{\gamma_i}{\sum_{j=1}^{M} \gamma_j}.$$

Let $I(\omega) \subset \{1, \ldots, M\}^{N(\omega)}$ be the set of simulated integers. Then we can simulate the total loss for the considered period as

$$X(\omega) := \sum_{i \in I(\omega)} X_i(\omega),$$

where the random variables (X_i) are independent and where for each i, X_i has distribution μ_i.
Said otherwise, we first simulate the number of claims, then we draw the events, and then the corresponding losses.

Aknowledgments The views expressed in the paper do not represent the views or position of CCR.

References

[1] Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. Coherent measures of risk. *Mathematical Finance 9*, 3 (1999), 203–228.

[2] Barrieu, P., and Scandolo, G. Assessing financial model risk. *European Journal of Operational Research 242*, 2 (2015), 546–556.

[3] Burden, R., and Faires, J. *Numerical Analysis.* Brooks/Cole, Cengage Learning, 2011.

[4] Chernozhukov, V., Fernández-Val, I., and Galichon, A. Quantile and probability curves without crossing. *Econometrica 78*, 3 (2010), 1093–1125.

[5] Clemen, R. T. Combining forecasts: A review and annotated bibliography. *International Journal of Forecasting 5* (1989), 559–583.

[6] de Vries, C. G., Samorodnitsky, G., Jorgensen, B. N., Mandira, S., and Danielsson, J. Subadditivity Re-Examined: the Case for Value-at-Risk. FMG Discussion Papers dp549, Financial Markets Group, Nov. 2005.

[7] Denis, L., and Martini, C. A theoretical framework for the pricing of contingent claims in the presence of model uncertainty. *Ann. Appl. Probab. 16*, 2 (05 2006), 827–852.

[8] Elliott, G., and Timmermann, A. *Economic Forecasting.* Princeton University Press, 2016.

[9] Fleming, W., and Soner, H. *Controlled Markov Processes and Viscosity Solutions.* Stochastic Modelling and Applied Probability. Springer New York, 2006.

[10] Genest, C., and Zidek, J. V. Combining probability distributions: A critique and an annotated bibliography. *Statist. Sci. 1*, 1 (02 1986), 114–135.

[11] Gneiting, T., and Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association 102*, 477 (2007), 359–378.

[12] Gneiting, T., and Ranjan, R. Combining predictive distributions. *Electron. J. Statist. 7* (2013), 1747–1782.

[13] Granger, C. W. J. Invited review combining forecasts twenty years later. *Journal of Forecasting 8*, 3 (1989), 167–173.

[14] Hou, Z., and Obłój, J. Robust pricing–hedging dualities in continuous time. *Finance and Stochastics 22*, 3 (Jul 2018), 511–567.

[15] Huber, P. *Robust Statistics.* Wiley Series in Probability and Statistics - Applied Probability and Statistics Section Series. Wiley, 2004.

[16] Hürlimann, W. Analytical bounds for two value-at-risk functionals. *ASTIN Bulletin 32*, 2 (2002), 235?265.

[17] McConway, K. J. Marginalization and linear opinion pools. *Journal of the American Statistical Association 76*, 374 (1981), 410–414.

[18] Nadaraya, E. Some new estimates for distribution functions. *Theory of Probability & Its Applications 9*, 3 (1964), 497–500.

[19] Soner, H. M., Touzi, N., and Zhang, J. Wellposedness of second order backward sdes. *Probability Theory and Related Fields 153*, 1 (Jun 2012), 149–190.

[20] Swiss Re Institute, . Sigma 2018. world insurance in 2017: Solid, but mature life markets weigh on growth.