December 2010 N°10

SCOR Papers

PrObEx:

A new method for the calibration of copula parameters from prior information,observations and expert opinions

By **PHILIPP ARBENZ**
ETH Zurich - SCOR Fellow

and **DAVIDE CANESTRARO**
SCOR SE

Abstract

A prudent assessment of dependence is crucial in many stochastic models in insurance mathematics. Copula models allow to represent dependence between random variables. However, copula estimation procedures usually contain a large parameter uncertainty if observations are scarce, which happens frequently in insurance applications. We propose a Bayesian method which combines prior information (e.g. from regulators), observations and expert opinions to estimate copula parameters and determine the estimation uncertainty. The combination of different sources of information can significantly reduce the parameter uncertainty compared to the use of only one source. We also describe methodology, psychological effects, and popular fallacies for obtaining expert opinions.

Keywords
Copulas, Expert judgment, Insurance, Dependence measure, Bayesian inference, Correlation, Risk management

SCOR Global Risk Center

1 Introduction

Insurance companies can model their liabilities by determining a stochastic model for both marginal distributions and dependence structure. Techniques for estimating marginal distributions from data and/or expert opinions are well known, see for instance McNeil et al. (2005) or Bühlmann and Gisler (2005). If many joint observations are available, estimating the dependence structure is relatively easy, see McNeil et al. (2005). In (re)insurance settings it frequently occurs that only very few *joint* observations are available, which may be the case even if plenty of information is available for the marginal distributions. In that case, it is often considered adequate to make the assumption of independence or impose simple assumptions on Pearson correlation coefficients. However, these approaches have been shown to contain several pitfalls when used in risk management, see Embrechts et al. (2002). From a regulatory point of view, the Solvency 2 guidelines enforce European insurance companies to assess dependence between their risks. See CEIOPS (2008) for the most recent report on the quantitative impact study on the new solvency guidelines. Similar requirements are imposed on those financial institutions which are subjected to the Basel II accords.

As the current financial crisis shows, dependence, particularly in the tails, must be accounted for correctly, see Donnelly and Embrechts (2010), Hellwig (2009) as well as Dacorogna and Canestraro (2010). Models for financial and insurance risks which account for dependence in a more comprehensive way than a variance-covariance approach (e.g. emphasizing joint extreme events) have gained much interest, see e.g. Denuit et al. (2005) or Embrechts et al. (2001).

One way to model dependence between random variables are the so called copula functions, which allow to separate the dependence structure from the margins. Copulas recently attracted much interest in finance and insurance, see McNeil et al. (2005) for their application in risk management and Genest et al. (2009) for an overview on developments in finance. In a symposium on insurance mathematics, Schnieper (2010) recommended to use copulas to avoid underestimating tail dependence.

Suppose an insurance company uses copulas to model dependence. If joint observations are scarce, the actuary may decide to also use other sources of information, such as expert judgment or regulatory guidelines, in order to find a good estimate of the copula parameters. For instance, certain joint extreme events might be anticipated by experts without having been observed yet. We are not aware of an existing sound mathematical framework to combine different sources of information in order to estimate copula parameters. This paper fills this gap by using a Bayesian framework within a parametric copula model and provides a method that is applicable even if observations are scarce.

We concentrate on the estimation of bivariate one-parameter copulas but our method can be extended to higher dimensions and more parameters. This work touches the fields of insurance mathematics, Bayesian statistics, psychology, and econometrics. For those readers who are not familiar with all of the four mentioned fields, we give extensive references to the relevant publications in the different research areas.

The paper is organized as follows. The problem setting is laid out in Section 2 and a Bayesian approach is described in Section 3. Issues surrounding the assessment of the prior density and the elicitation of expert opinions are studied in Section 4 and Section 5, respectively. A simple application is discussed in Section 6, after which we conclude in Section 7.

2 Copulas and Dependence Measures

Let $(X, Y) \in \mathbb{R}^2$ be a bivariate random vector on some probability space $(\Omega, \mathfrak{A}, \mathbb{P})$. Assume the margins $F(x) = \mathbb{P}[X \leq x]$ and $G(y) = \mathbb{P}[Y \leq y]$ of (X, Y) are continuous. The copula representation of the joint cumulative distribution function (cdf) $H(x, y) = \mathbb{P}[X \leq x, Y \leq y]$ is given by

$H(x, y) = C\left(F(x), G(y)\right)$, for all $x, y \in \mathbb{R}$,

where $C : [0, 1]^2 \to [0, 1]$ is the unique copula function, see Theorem 2.3.3 in Nelsen (2006). The copula C is also the cdf of the random vector $(U, V) = (F(X), G(Y)) \in [0, 1]^2$, denoted with $(U, V) \sim C$. Note that (U, V) has uniform margins. We refer to Nelsen (2006) for a detailed introduction to copulas. From a practical point of view, it is often convenient to model margins and the copula separately, with different parametric families. Four popular parametric copula families are given in Appendix A.

Let $\mathcal{C}_0 = \{C_\gamma : \gamma \in \Gamma\}$ denote a family of bivariate differentiable copulas, with parameter set Γ and density $c(\cdot|\gamma)$. For the rest of the paper, we assume that

$$C = C_{\gamma_0} \in \mathcal{C}_0$$

where γ_0 is an unknown but fixed parameter. Our aim is to estimate γ_0.

In order to measure the degree of dependence between X and Y (or equivalently, between U and V), different dependence measures can be used, for example Spearman's rho, Kendall's tau or tail dependence. Let $\rho(\cdot, \cdot)$ denote a fixed dependence measure. Define the set of attainable values of ρ for copulas in $\mathcal{C}_0$ by

$$\Theta = \{\rho(U^*, V^*) : (U^*, V^*) \sim C^* \in \mathcal{C}_0\}.$$

We assume that Θ is an interval, i.e. $\Theta = [a, b] \subset \mathbb{R}$ and justify this assumption later. Five well known dependence measures are given in Appendix B, which satisfy $\rho(U, V) = \rho(X, Y)$, i.e. are not dependent on the margins. As dependence measures are more intuitive for experts to assess (details are given in Section 5), we parameterize the copula family $\mathcal{C}_0$ in terms of the dependence measure. In the sequel, we assume there exists a bijective link function $g : [a, b] \to \Gamma$ such that

$g\left(\rho(U^*, V^*)\right) = \gamma^*$
for all (U^*, V^*) with $(U^*, V^*) \sim C_{\gamma^*} \in \mathcal{C}_0$,

i.e. each value $\theta^* \in [a, b]$ can be mapped to a unique copula parameter $\gamma^* = g(\theta^*) \in \Gamma$ such that $\rho(U^*, V^*) = \theta^*$ for $(U^*, V^*) \sim C_{\gamma^*} = C_{g(\theta^*)}$. Let

$$\theta_0 = \rho(U, V),$$

which is fixed but unknown and note that θ_0 satisfies $g(\theta_0) = \gamma_0$. In Appendix C, we show that for most combinations of copula families shown in Appendix A and dependence measures shown in Appendix B, both the assumptions on Θ being an interval and the existence of g hold.

In the following, we describe the statistical problem under study. Instead of estimating γ_0 directly, we calculate an estimate $\widehat{\theta_0}$ of θ_0, which leads to an estimate $\widehat{\gamma_0} = g(\widehat{\theta_0})$ of γ_0. We assume that (up to) three sources of information are available:

1. A prior source of information, e.g. regulator guidelines, to which a prior density $\pi(\theta) : [a, b] \to [0, \infty)$ can be fitted. More details will be given in Section 4.
2. N independent observations (U_n, V_n), $n = 1, \ldots, N$, of $(U, V) \sim C_{\gamma_0}$. The set of observations is denoted by $\mathcal{O} = \{(U_n, V_n) : n = 1, \ldots, N\}$.
3. K experts, each providing one point estimate φ_k, $k = 1, \ldots, K$, of $\rho(U, V)$. The set of expert assessments is denoted by $\mathcal{E} = \{\varphi_k : k = 1, \ldots, K\}$.

In Section 3, we will introduce the Bayesian method which allows to estimate γ_0 based on the three mentioned sources.

If F and G are unknown, the true observations $\{(X_n, Y_n) : n = 1, \ldots, N\}$ of $(X, Y) \sim H$ can be transformed to pseudo-observations of $(U, V) \sim C$ by setting $(U_n, V_n) = (\widehat{F}(X_n), \widehat{G}(Y_n))$, $n = 1, \ldots, N$, where $\widehat{F}$ and $\widehat{G}$ are estimations of F and G. These estimations can be obtained parametrically or non-parametrically. One possibility are the modified empirical cdf $F(x) = \frac{1}{N+1}\#\{n : X_n \leq x\}$ and $\widehat{G}(x) = \frac{1}{N+1}\#\{n : Y_n \leq x\}$.

When using $\widehat{F}$ and $\widehat{G}$ instead of F and G, the resulting (U_n, V_n) are not strictly independent. However, as the aim of this paper is to focus on estimating dependence, we assume the margins to be known.

3 The Use of Bayesian Inference to Combine Different Sources of Information

In this section, we present the Bayesian model that is used to combine the three sources of information mentioned above. For a detailed introduction to Bayesian inference, see Lindley (1965) or Bernardo and Smith (1994). Our method is based on Lambrigger et al. (2007) who apply Bayesian inference to combine three sources of information in order to estimate regulatory capital for operational risks.

We will first infer a Bayesian point estimate of θ_0, which can then be mapped to an estimate of γ_0. Let the prior density $\pi(\theta) : [a, b] \to [0, \infty)$ represent the initial (subjective) information available on θ_0. A methodology for determining $\pi(\theta)$ is given in Section 4. From a frequentist perspective, θ_0 can be seen as a realization of a random variable θ with density $\pi(\theta)$.

The additional information given by $\mathcal{O}$ and $\mathcal{E}$ is used to adapt $\pi(\theta)$. That is, we replace the prior density $\pi(\theta)$ by a posterior density $\pi(\theta|\mathcal{O}, \mathcal{E})$ of θ given $\mathcal{O}$ and $\mathcal{E}$. Bayes' Theorem leads to the relation

$$\pi(\theta|\mathcal{O}, \mathcal{E})h(\mathcal{O}, \mathcal{E}) = h(\mathcal{O}, \mathcal{E}|\theta)\pi(\theta), \qquad (1)$$

where $h(\mathcal{O}, \mathcal{E})$ denotes the unconditional and $h(\mathcal{O}, \mathcal{E}|\theta)$ the conditional, given θ, joint density of $\mathcal{O}$ and $\mathcal{E}$. We will, under suitable assumptions, express $h(\mathcal{O}, \mathcal{E}|\theta)$ in terms of more elementary functions in order to make the posterior $\pi(\theta|\mathcal{O}, \mathcal{E})$ numerically accessible. Similar to Lambrigger et al. (2007) we assume that the expert assessments and the observations are independent. It follows that

$$h(\mathcal{O}, \mathcal{E}|\theta) = h_O(\mathcal{O}|\theta)h_E(\mathcal{E}|\theta), \qquad (2)$$

where h_O and h_E are the conditional densities, given θ, of $\mathcal{O}$ and $\mathcal{E}$, respectively.

As the observations (U_n, V_n), $n = 1, \ldots, N$, are independent,

$$h_O(\mathcal{O}|\theta) = \prod_{n=1}^{N} c\left(U_n, V_n|g(\theta)\right), \qquad (3)$$

where $c(u, v|\gamma) = \frac{\partial}{\partial u}\frac{\partial}{\partial v}C_\gamma(u, v)$.

Assuming the experts form their opinions independently of each other, we can write

$$h_E(\mathcal{E}|\theta) = \prod_{k=1}^{K} e_k(\varphi_k|\theta), \qquad (4)$$

where $e_k(\cdot|\theta)$ is the conditional density, given θ, of the k-th expert assessment. We model $\varphi_k|\theta$ with a shifted Beta distribution on $[a, b]$ with mean $\mathbb{E}[\varphi_k|\theta] = \theta$ (i.e. conditionally unbiased) and var$(\varphi_k|\theta) = \sigma_k^2$. A discussion of the assumptions on $\varphi_k|\theta$ as well as methods to estimate σ_k^2 are given in Section 5.3 and Appendix D.

Note that $h(\mathcal{O}, \mathcal{E})$ does not depend on θ. Hence, plugging (2), (3), and (4) into (1), yields

$$\pi(\theta|\mathcal{O}, \mathcal{E}) \propto \pi(\theta) h(\mathcal{O}, \mathcal{E}|\theta)$$
$$= \pi(\theta) \prod_{n=1}^{N} c(U_n, V_n|g(\theta)) \prod_{k=1}^{K} e_k(\varphi_k|\theta),$$

where the symbol $\propto$ denotes proportionality with respect to θ.

For sensitivity analysis or in case no expert opinions or observations are available, we can also compute the posterior of θ given only $\mathcal{O}$ or $\mathcal{E}$:

$$\pi(\theta|\mathcal{O}) \propto \pi(\theta) \prod_{n=1}^{N} c(U_n, V_n|g(\theta)),$$
$$\pi(\theta|\mathcal{E}) \propto \pi(\theta) \prod_{k=1}^{K} e_k(\varphi_k|\theta).$$

Any Bayesian point estimator (such as mean, mode or median) can be used to calculate a point estimate $\widehat{\theta_0}$ from $\pi(\theta|\mathcal{O}, \mathcal{E})$. We will use the posterior mean $\widehat{\theta_0} = \mathbb{E}[\theta|\mathcal{O}, \mathcal{E}]$ in order to be consistent with the modeling of the expert assessments, which is based on matching conditional moments. The uncertainty of $\widehat{\theta_0}$ can be assessed through var$(\theta|\mathcal{O}, \mathcal{E})$.

We conclude this section with several remarks.

In other applications of Bayesian inference, conjugate priors are often used, in which case both $\pi(\theta)$ and $\pi(\theta|\mathcal{O}, \mathcal{E})$ belong to the same parametric class of distributions, see Bernardo and Smith (1994). We are not aware of any conjugate priors for copulas.

Although the influence of the experts on each other may be reduced by minimizing communication during the elicitation, experts can be dependent as they can have the same basis of knowledge. If the independence assumption for experts is deemed too strong, dependent experts can be modeled as shown in Jouini and Clemen (1996). However, according to Kallen and Cooke (2002), experts are in general less dependent than anticipated.

Dependence between experts and observations is difficult to avoid, but it is unclear how this influences the expert assessments. An expert might tend to underestimate due to inexisting joint observations or to overestimate by overcorrecting through unrealistic scenarios.

In most cases the normalization factor $h(\mathcal{O}, \mathcal{E})$ is not known. Since $\Theta = [a, b]$ is an interval it is easy to find an approximation to $\pi(\theta|\mathcal{O}, \mathcal{E})$: the interval $[a, b]$ can be discretized with a grid on which $\pi(\theta)h(\mathcal{O}, \mathcal{E}|\theta)$ can be calculated and rescaling yields an approximation of $\pi(\theta|\mathcal{O}, \mathcal{E})$. If a grid approximation is not feasible, Markov Chain Monte Carlo (MCMC) methods can be used, see Robert and Casella (2005).

As we intend our method to be used mainly with small N and K, we refrain from proving asymptotic statements on the convergence of $\theta|\mathcal{O}, \mathcal{E}$ to θ_0. Proofs and more details concerning the following statements can be found in Section 10 of Van der Vaart (1998). For $N \to \infty$ the information contained in $\mathcal{O}$ increases, the influence of $\pi(\theta)h_E(\mathcal{E}|\theta)$ is diminished, and the posterior is driven by the observations. The Bernstein-von Mises Theorem states asymptotic normality of $\theta|\mathcal{O}, \mathcal{E}$ for $N \to \infty$ if $\pi(\theta)h_E(\mathcal{E}|\theta)$ is smooth and positive in a neighborhood of θ_0. Therefore, $\theta|\mathcal{O}, \mathcal{E}$ and thus also $\widehat{\theta_0}$ converge in probability to θ_0 with rate $1/\sqrt{N}$. Furthermore, Bayesian point estimators are asymptotically efficient and asymptotically equivalent to maximum likelihood estimators. If both prior and expert assessments are modeled with a shifted Beta, the conditions of the Bernstein-von Mises Theorem are satisfied.

4 Assessing the Prior Distribution

The value $\pi(\theta)$ reflects the prior likelihood that $\theta = \theta_0$, or equivalently, that $C_{g(\theta)} = C_{\gamma_0}$. In Bayesian terms, $\pi(\theta)$ is the unconditional density of the random variable θ.

Suppose we can infer a point estimate $\widehat{\theta}_p$ of θ_0 (p for prior) from the prior source of information, e.g. regulator guidelines. We then propose to model $\pi(\theta)$ with a shifted Beta distribution with mean $\mathbb{E}[\theta] = \widehat{\theta}_p$. The variance var$(\theta)$ determines the credibility which is given to $\widehat{\theta}_p$. If the source of information leading to $\widehat{\theta}_p$ does not specify a measure of uncertainty, we propose to determine var(θ) through a qualitative approach according to the subjective confidence that is given to the estimate $\widehat{\theta}_p$:

$$\text{var}(\theta) = \begin{cases} 0.005(b-a)^2 & \text{for high confidence,} \\ 0.02(b-a)^2 & \text{for intermediate confidence,} \\ 0.05(b-a)^2 & \text{for low confidence.} \end{cases}$$

Though this approach may seem arbitrary, we found it to be practical in the sense that it allows the statistician to easily assign a credibility to the prior estimate $\widehat{\theta}_p$.

If no prior belief is available then $\pi(\theta)$ can be set uninformative, i.e. such that the posterior distribution depends mainly on $\mathcal{O}$ and $\mathcal{E}$, see Price and Manson (2002) for an introduction to uninformative priors and Berger and Bernardo (1992) for a critical view. For Θ being an interval, an uninformative prior is given by the uniform distribution on $[a, b]$.

Figure 1 illustrates the four mentioned qualitative approaches.

The following list gives three possibilities for the prior source of information.

1. **Regulator guidelines.** Some insurance regulators publish reference values for the correlation between certain risk types. See for instance Section SCR.9 in CEIOPS (2010) and Section 10.1.4.5 in CEIOPS (2008) for the proposals of the European Union regulators as well as Section 8.4 in BPV (2006) for directives from the Swiss regulators.

2. **Physically similar situations.** Analogous to the proposal in Lambrigger et al. (2007), $\widehat{\theta}_p$ can be taken as the known degree of dependence of two random variables whose dependence is similar in nature to the dependence between X and Y. For instance, the dependence between fire losses in one city might be similar to those in a neighbouring city. This is similar to the approaches used in credibility theory where collective data can be used as a starting point to estimate individual parameters.

3. **Expert Judgement.** An expert can either estimate $\widehat{\theta}_p$ according to his personal belief or in order to incorporate an artificial bias, e.g. by putting weight on high degrees of dependence in order to avoid underestimating dependence.

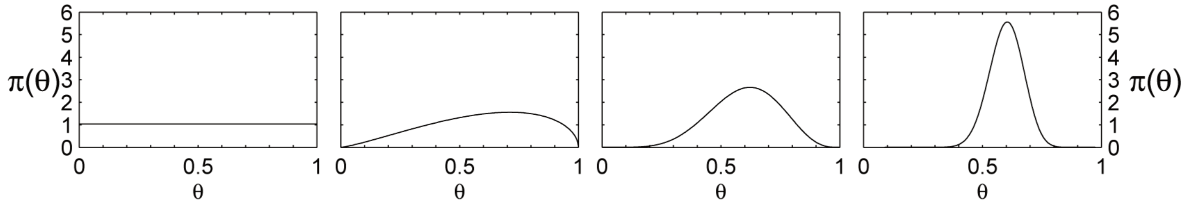

Figure 1: Four Beta densities on the interval $[0, 1]$. First: uniform and uninformative. Second to fourth: $\widehat{\theta}_p = \mathbb{E}[\theta] = 0.6$ with low, intermediate and high confidence, respectively. The corresponding Beta parameters are: $(\alpha, \beta) = (0, 0)$; $(2.28, 1.52)$; $(6.6, 4.4)$; $(28.2, 18.8)$.

5 The Elicitation and Modeling of Expert Opinion

As observations are lacking or sparse in many situations in the context of risk management, expert opinion is increasingly recognized as an additional source of relevant information, see Cooke and Goossens (2000). The Swiss regulator of insurances explicitly considers expert judgment for model parameter estimation, see FINMA (2008). A useful annotated bibliography focussing on expert judgment is provided by Genest and Zidek (1986). For an overview of the more recent literature on the use of expert opinion see Garthwaite et al. (2005) and Ouchi (2004).

A person is called an expert if he is able to give mathematically substantial answers in the area of interest. A procedure to find adequate experts is given in Cooke and Goossens (2000).

According to Clemen and Winkler (1999) there are two scientifically advocated approaches to attain and combine expert opinion. In *behavioral approaches*, experts interact and agree on a common conclusion by means of discussions and other forms of interaction. Two examples of behavioral approaches are given in Delbecq et al. (1975) and Linstone and Turoff (1975). In *mathematical approaches*, experts are elicited separately. Then their opinions are expressed as subjective probabilities and subsequently combined through mathematical methods.

According to Mosleh et al. (1988), behavioral approaches have the disadvantage that they are prone to be influenced by dominant personalities, they can suffer from the limited participation of less confident experts, and there is a general tendency to reach a conclusion too fast. Daneshkhah (2004) concludes that mathematical approaches mostly provide more reliable results than behavioral approaches.

The task of assessing dependence has received little attention. The approach to directly elicit the value of a dependence measure has been subjected to criticism as well as appraisal, see Morgan and Henrion (1992) and Clemen et al. (2000), respectively.

Our approach of asking experts for estimates of θ_0 and applying Bayesian inference represents a mathematical approach. Note that other mathematical approaches allow experts to formulate their answer in the form of paired comparisons, confidence intervals or whole distributions, see Meyer and Booker (2001) for an overview. We focus on point estimates as this approach allows a direct representation in a Bayesian framework.

Turning expert opinions into scientifically meaningful statements requires adherence to psychological and procedural principles, given in Section 5.1 and 5.2, respectively. Section 5.3 discusses the modeling of expert assessments and the estimations of their uncertainty.

5.1 Psychological Effects in Expert Elicitation

In order to comprehensively understand a process involving expert judgment, one must understand the psychological effects involved when experts are asked to quantify their opinion. Different approaches have been proposed to understand and categorize these effects, see Meyer and Booker (2001).

Three categories are used to describe the consistency of expert assessments.

1. An expert is *coherent* if his probability assessments collectively do not violate formal probability theory.
2. An expert is *calibrated* if, on average, his opinions correspond to the (eventual) true outcomes.
3. An expert is *reliable* if he agrees with himself in (theoretically) reiterated elicitations.

Most experts are not familiar with describing their beliefs in terms of probability assessments. However, coherence can be achieved through appropriate questioning and training in probability calculus, see Kynn (2005). Providing feedback to the experts helps them to become well calibrated on their specific task, see Lichtenstein and Fischhoff (1980). Wallsten and Budescu (1983) find that experts are sufficiently reliable when questioned within their area of expertise.

The so called *cognitive models*, describe the processes that are involved in assessing probabilities and investigate which kind of distortions will affect expert assessments. An overview on cognitive models in psychology literature is given in Kynn (2008).

In contrast to the cognitive models, the *heuristics and biases* research stream identified different instinctive strategies to judge under uncertainty, which influence expert assessments, see Tversky and Kahneman (1974). These strategies were observed to produce quick decisions but as well to induce certain biases, see O'Hagan et al. (2006). The heuristics and biases research program was debated and severely criticized in the psychological research, see Kynn (2008). However, recent statistical literature seems to ignore this critique to a large extent and a general advice is to mention such biases to the experts before the elicitation begins, see Rebonato (2010).

In the following, we describe some examples of involved psychological effects.

- In a study described in Kahneman and Tversky (1982), people were asked: "Linda is 31 years old, single, outspoken, bright and majored in philosophy. She is deeply concerned with issues of discrimination and social justice. Which is more likely? (i) Linda is a bank teller; (ii) Linda is a bank teller who is active in the feminist movement." Most answers rank (ii) to be more probable than (i), not considering that (ii) must have a probability less or equal than (i) because (ii) is a subset of (i).
- According to Eddy (1982), doctors are likely to confuse $\mathbb{P}(\text{positive test}|\text{disease})$ (i.e. the test sensitivity) with $\mathbb{P}(\text{disease}|\text{positive test})$ (i.e. the power of the test).
- People tend to ignore prior probabilities when judging a conditional probability. When asked "Is a meticulous, introverted, meek and solemn person more likely to be engaged as a librarian or as a salesman?" people mostly answer librarian as the stereotype of a librarian better suits the characteristics of this person, see Kahneman and Tversky (1973). However, this answer ignores the fact that there are much more salesmen than librarians.
- The probability assigned to an event increases with the amount of details that are given to describe it. For instance, the estimated probability that a person dies due to a natural cause is usually smaller than the sum of the separately estimated probabilities for heart disease, diabetes and other natural causes, see O'Hagan et al. (2006).

5.2 The Elicitation Process

The general goal of applying expert elicitation procedures is to allow decisions to be taken in a rational manner and to be perceived to be as such. For this reason, an expert elicitation procedure must be subjected to five principles, see Cooke (1991):

1. **Reproducibility.** All data (in particular questionnaires and the elicited opinions) must be open to qualified reviewers and results must be reproducible in order to allow revision from auditors or regulators.
2. **Accountability.** Filled questionnaires are stored and each opinion can be linked to the corresponding expert.
3. **Empirical control.** There should be in principle (at least in the far future) the possibility to verify expert opinions on the basis of measurable observations.

4. **Neutrality.** The process must induce experts to honestly state their opinion. There must not be any incentives (such as changes in salary or reputation) for the experts to give answers different from their true opinion.
5. **Fairness.** Experts are not discriminated or given smaller weights if not due to reasons that can be justified within the mathematical approach.

An estimation of the parameter γ_0 constitutes the final goal of the study. It is inconvenient to directly elicit γ_0 as it is not an observable quantity and may not be familiar to the expert in terms of the way he collects and evokes his knowledge. However, a question asking for the value of a dependence measure can be formulated in such a way that substantial answers can be given even if experts are unfamiliar with probability theory. For this reason, we ask experts to estimate θ_0.

The expert elicitation process can generally be split into three parts, where each requires adherence to certain guidelines, see Cooke and Goossens (2000):

- During *pre-elicitation*, the questionnaire is written and tested through a dry-run. A well developed questionnaire can significantly improve the consistency of the expert assessments. A training session must be conducted, where heuristics, formal probability theory, and the meaning of the questions are addressed, see Clemen et al. (2000).
- In the *elicitation*, each expert individually quantifies his opinion. A person who can resolve ambiguities in the procedure or in the questionnaire should be present.
- The *post-elicitation* covers the processing of the expert assessments. A sensitivity analysis is recommended, where the importance of model assumptions as well as the impact of each expert on the result is investigated.

5.3 Bayesian Modeling and Assessing Expert Uncertainty

We will now address the mathematical modeling of the expert assessments $\mathcal{E}$. Adhering to Bayesian principles, the estimates φ_k are assumed to be realizations of random variables. There is no canonical way to model the distribution of these random variables, hence we will construct a model consistent with our Bayesian framework.

As we believe our experts to be calibrated, we model the expert estimates to be conditionally unbiased, i.e. $\mathbb{E}[\varphi_k|\theta] = \theta$ for all $\theta \in [a, b]$, as it is also used in Lambrigger et al. (2007). To reflect the expert uncertainty we assign each expert a variance σ_k^2, $k = 1, \ldots, K$, which is assumed to be independent of θ: $\text{var}(\varphi_k|\theta) = \sigma_k^2$ for all $\theta \in [a, b]$. Note that a weighting of the expert estimates through their variances determines the amount of information that is contained in the corresponding estimate, i.e. we do not have to further specify the relative weights between the observations and the experts. We will illustrate below how to estimate σ_k^2.

From a modeling perspective, the only remaining task is to fit a conditional density which attains these moments. Due to the versatility and the explicit expressions for moments and because Θ is assumed to be an interval we use the (shifted) Beta distribution. Other distributions on intervals such as the Kumaraswamy, triangular and raised cosine distributions were tested. However, these have difficult expressions for moments or are not smooth with a bounded support, hence they were found to be inferior to the Beta distribution from our Bayesian modelling viewpoint.

We show three possible approaches to calculate estimates $\widehat{\sigma_k^2}$ of σ_k^2, where the most suitable approach must be chosen according to the available information.

1. **Homogeneous experts.** If experts are assumed to have equal uncertainty, i.e. $\sigma_k^2 = \sigma^2$ for $k = 1, \ldots, K$, we may estimate

$$\widehat{\sigma^2} = \frac{1}{K-1}\sum_{k=1}^{K}\left(\varphi_k - \overline{\varphi}\right)^2, \quad \text{where:}$$

$\overline{\varphi} = \frac{1}{K}\sum_{k=1}^{K}\varphi_k$.

2. **Seed variables.** Suppose we have a number of H seed variables. Seed variables are values $\psi_0^{(h)}$, $h = 1, \ldots, H$, which are known to the person doing the elicitation but not known to the experts. The experts are then asked to estimate the seed variables. By assuming that the experts are unbiased and that their uncertainty in estimating the seed variables is the same as the uncertainty when estimating θ_0, we can estimate σ_k^2 from the differences $\psi_k^{(h)} - \psi_0^{(h)}$. In mathematical terms, we suppose that $\mathbb{E}[\psi_k^{(h)}|\psi_0^{(h)}] = \psi_0^{(h)}$ and $\text{var}(\psi_k^{(h)}|\psi_0^{(h)}) = \sigma_k^2$ for $h = 1, \ldots, H$ and $k = 1, \ldots, K$. We can then estimate σ_k^2 with

$$\widehat{\sigma_k^2} = \frac{1}{H}\sum_{h=1}^{H}\left(\psi_k^{(h)} - \psi_0^{(h)}\right)^2, \quad \text{for } k = 1, \ldots, K.$$

Seed variables should not be general knowledge variables but they have to belong to the experts' field of expertise or to an adjacent field, see Cooke et al. (1988). Ideal seed variables for the purpose of our model are given by the value of the dependence measure ρ applied to random vectors which also lie in the field of expertise of the expert. If H is low, the $\psi_k^{(h)}$ can be used to estimate an equal variance for all experts, similar to approach 1.

3. **Subjective variances.** The σ_k^2 can be estimated through any technique deemed feasible, e.g. through the number of years of the experts' experience or through the subjective judgment of another expert, see Gokhale and Press (1982). This includes a weighted average between approaches 1, 2 and 3. Winkler (1968) proposes to let experts provide an estimate of their uncertainty themselves, which is however criticized in Cooke et al. (1988), as experts tend to be overconfident in their estimates.

6 An Application to Risk Management

In this section, we give an illustration of the above ideas, showing the use of our method to estimate the parameter of a t-copula. Furthermore, we apply the method to estimate the value and to quantify the estimation uncertainty of the Value-at-Risk (VaR) of the sum of two random variables.

Let $\mathcal{C}_0 = \{C_{3,\gamma}^t : \gamma \in [-1,1]\}$, i.e. we aim to estimate the parameter γ_0 of a t-copula with three degrees of freedom (dof). Let the dependence measure be Kendall's tau, $\rho = \rho_\tau$, which implies $\Theta = [a,b] = [-1,1]$. As $\rho_\tau = 2/\pi \arcsin(\gamma)$ for t-copulas, the link function is $g(\theta) = \sin(\theta\pi/2)$, see McNeil et al. (2005). Suppose we have no information available to fit the prior to, hence we use the uniform distribution $\pi(\theta) = 1/2$ on $[-1,1]$ as an uninformative prior.

Let $N = 24$ copula observations be given as: $\mathcal{O} = \{$ (0.12,0.04), (0.04,0.12), (0.08,0.16), (0.20,0.24), (0.48,0.08), (0.36,0.40), (0.60,0.20), (0.16,0.68), (0.24,0.64), (0.40,0.48), (0.32,0.56), (0.28,0.72), (0.76,0.28), (0.52,0.60), (0.72,0.44), (0.44,0.76), (0.88,0.36), (0.96,0.32), (0.80,0.52), (0.64,0.80), (0.56,0.92), (0.68,0.96), (0.84,0.88), (0.92,0.84) $\}$, which are illustrated in Figure 2.

Assume three experts ($K = 3$) gave estimates $\varphi_1 = 0.35$, $\varphi_2 = 0.5$, and $\varphi_3 = 0.7$ of the Kendall's tau $\rho_\tau(U,V)$. We model the expert estimates as described in Section 5.3 with a shifted Beta

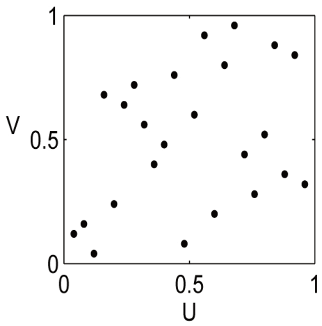

Figure 2: A scatterplot of the observations $\mathcal{O}$.

	Seed variables				
True value	0.3	0.7	0.5	0.55	
	Expert assessments				Estimated variance
Expert 1	0.1	0.3	0.6	0.75	0.0625
Expert 2	0.5	0.5	0.35	0.6	0.0262
Expert 3	0.3	0.8	0.3	0.8	0.0281

Table 1: The estimates from the three experts ($k = 1, 2, 3$) of the four seed variables ($h = 1, \ldots, 4$) with true value $\psi_0^{(h)}$. These estimates are used to calculate $\widehat{\sigma_1^2}$, $\widehat{\sigma_2^2}$, and $\widehat{\sigma_3^2}$.

distribution with endpoints -1 and $+1$.

To estimate the experts' variances σ_k^2, we use the approach (ii) using four seed variables as shown in Table 1. This leads to the variance estimates $\widehat{\sigma_1^2} = 0.0625$, $\widehat{\sigma_2^2} = 0.0262$, and $\widehat{\sigma_3^2} = 0.0281$.

The densities $\pi(\theta|\mathcal{O})$, $\pi(\theta|\mathcal{E})$, and $\pi(\theta|\mathcal{O}, \mathcal{E})$ can now be computed numerically and are shown along with $\pi(\theta)$ in Figure 3. The point estimates, standard deviations and 90% credible intervals resulting from those distributions are given in Table 2.

The best estimate $\widehat{\gamma_0}$ using all information is then

$$\widehat{\gamma_0} = g(\mathbb{E}[\theta|\mathcal{O}, \mathcal{E}]) = g(0.399) = 0.587.$$

Suppose that we are interested in the unknown but fixed quantity $\text{VaR}_{0.99}(X + Y|\theta_0)$, where both $X = F^{-1}(U)$ and $Y = G^{-1}(V)$ have a t-margin with three dof and $\mathbb{E}[X] = 1$, $\mathbb{E}[Y] = 2$, $\text{var}(X) = 192$ and $\text{var}(Y) = 243$.

This implies that for a given copula parameter γ, we have that (X, Y) is bivariate-t with three dof and:

$$\mathbb{E}[(X, Y)|\gamma] = (1, 2),$$

$$\text{Cov}(X, Y|\gamma) = \begin{pmatrix} 192 & 216\gamma \\ 216\gamma & 243 \end{pmatrix}.$$

As the bivariate-t distribution is an elliptical distribution, we have for a given γ that $(X + Y)|\gamma$ is t-distributed with $\mathbb{E}[X + Y|\,\gamma] = 3$ and $\text{var}(X+Y|\gamma) = 435+432\gamma$. Because $\gamma = \sin(\theta\pi/2)$, we have

$$\begin{aligned}&\text{VaR}_{0.99}(X + Y|\theta)\\ &= 3 + \sqrt{145 + 144 \sin(\theta\pi/2)}\, q_3^{-1}(0.99),\end{aligned}$$

where q_3^{-1} is the inverse df of a standard-t distribution with three dof and $q_3^{-1}(0.99) \approx 4.5407$.

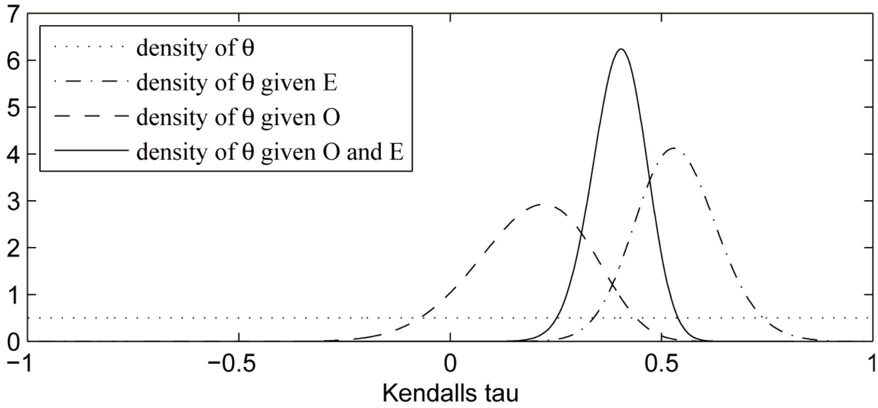

Figure 3: The densities $\pi(\theta)$, $\pi(\theta|\mathcal{O})$, $\pi(\theta|\mathcal{E})$ and $\pi(\theta|\mathcal{O},\mathcal{E})$.

	$\mathbb{E}[\cdot]$	$\sqrt{\text{var}(\cdot)}$	90% credible interval
θ	0.000	0.577	$[-0.900,\ \ 0.900]$
$\theta\|\mathcal{E}$	0.537	0.098	$[\ \ 0.380,\ \ 0.703]$
$\theta\|\mathcal{O}$	0.192	0.135	$[-0.046,\ \ 0.399]$
$\theta\|\mathcal{O},\mathcal{E}$	0.399	0.064	$[\ \ 0.291,\ \ 0.502]$

Table 2: Point estimates of $\rho_\tau(U,V)$ and associated uncertainties, measured through standard deviation and a credible interval.

The minimally and maximally attainable values for $\text{VaR}_{0.99}(X+Y|\theta)$ are 7.54 and 80.2, which correspond to $\theta=\gamma=\pm1$.

The densities of $\text{VaR}_{0.99}(X+Y)$, $\text{VaR}_{0.99}(X+Y|\mathcal{O})$, $\text{VaR}_{0.99}(X+Y|\mathcal{E})$, and $\text{VaR}_{0.99}(X+Y|\mathcal{O},\mathcal{E})$ are given in Figure 4 and the implied estimates of $\text{VaR}_{0.99}(X+Y|\theta_0)$ are given in Table 3.

This example shows how the uncertainty in estimating the dependence carries over to the uncertainty in the risk measure. Figure 4, illustrates that adding information ($\mathcal{O}$ and/or $\mathcal{E}$) reduces the estimation uncertainty.

In practice, quantiles are mostly estimated through parametric models (estimated with maximum likelihood) or through techniques stemming from extreme value theory, such as the Peaks-over-Threshold method (POT), see Embrechts et al. (2001). However, these techniques do not allow to incorporate expert opinions and confidence intervals derived from them may be of dubious quality for a small sample size. On the other hand, Bayesian inference allows a natural representation of parameter uncertainty. Modeling also the marginal distributions of X and Y in a Bayesian framework would be a direct extension of our method.

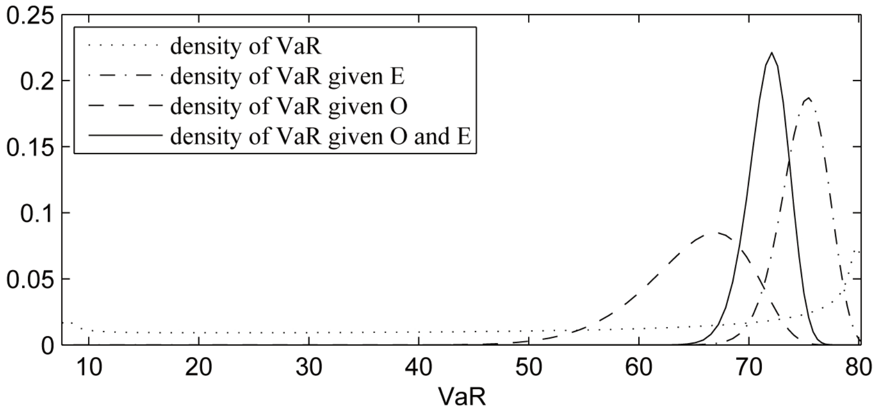

Figure 4: The densities of VaR$_{0.99}(X+Y)$, VaR$_{0.99}(X+Y|\mathcal{O})$, VaR$_{0.99}(X+Y|\mathcal{E})$ and VaR$_{0.99}(X+Y|\mathcal{O},\mathcal{E})$.

	$\mathbb{E}[\cdot]$	$\sqrt{\text{var}(\cdot)}$	90% credible interval
VaR$_{0.99}(X+Y)$	65.35	21.16	$[13.67, 80.17]$
VaR$_{0.99}(X+Y|\mathcal{E})$	75.47	2.181	$[71.65, 78.82]$
VaR$_{0.99}(X+Y|\mathcal{O})$	65.57	4.844	$[56.58, 72.29]$
VaR$_{0.99}(X+Y|\mathcal{O},\mathcal{E})$	71.89	1.827	$[68.69, 74.67]$

Table 3: Point estimates of VaR$_{0.99}(X+Y|\theta_0)$ and associated uncertainties, measured through standard deviation and a credible interval.

7 Conclusion

Based on Bayesian inference, we proposed a method to estimate a copula parameter by combining three sources of information, namely prior information, observations, and expert opinions. Our model allows to compute estimates of copula parameters and assess their uncertainty. The uncertainty and estimates of functionals of a copula (e.g. risk measures) can also be assessed. The model can also be used if not all of the three sources of information are available. For instance, the model allows to estimate a copula parameter from expert opinion only ($N = 0$) or from observations only ($K = 0$).

Our method is most helpful in situations where observations are scarce, i.e. in cases when concurring methods like maximum-likelihood usually exhibit severe parameter uncertainties. Incorporating different sources of information then strongly reduces the parameter uncertainty and parameter risk. The method converges asymptotically normal, i.e. as fast as maximum likelihood methods.

We investigated the challenging process of turning expert opinions into quantitative information. The Bayesian approach allows a natural interpretation of the expert assessments. Certain principles deduced from psychological and statistical research must be adhered to in order to get usable results. We proposed procedures to assess the accuracy of the expert assessments through estimating their variance, which controls their weight in the final estimate.

The extension of the method to high-dimensional copulas with only one parameter, such as the Clayton copula, is straightforward. However, modeling copulas with more parameters is an interesting topic for future research. Investigating the impact of dependent experts as well as the possibility of also modeling margins in a Bayesian framework promise to be interesting.

Acknowledgements

The authors would like to thank Michel Dacorogna and Paul Embrechts for their advice and insightful comments. Furthermore, Hans Bühlmann, Raffaele Dell'Amore, Christian Genest, Marius Hofert, Dominik Lambrigger and Johanna Nešlehová provided valuable remarks.

The acronym "PrObEx" (which stands for prior, observations and experts) mentioned in the title has been kindly suggested to us by Michel Dacorogna. The first named author thanks SCOR SE for financial support.

Appendix

A Four Popular Copula Families

In Table 4, we give the definitions of four popular bivariate copula families and their parameter set Γ. An extensive list of parametric copula families is given in Nelsen (2006).

If the copula family is not determined, Bayesian methods can be used to choose the best copula family among a set of families, see Huard et al. (2006).The hypothesis that the copula belongs to a specific family for a set of observations can be tested with goodness-of-fit tests as shown in Beaudoin et al. (2009).

It is difficult to set a prior for the dof ν of the t-copula in a senseful manner, hence we propose to set $\nu > 0$ to be a fixed number.

Copula family	Definition	Parameter range Γ
Gaussian	$C_\gamma^{Ga}(u,v) = \Phi_\gamma\left(\varPhi^{-1}(u), \varPhi^{-1}(v)\right)$	$\gamma \in (-1,1)$
t	$C_{\nu,\gamma}^{t}(u,v) = \mathbf{t}_{\nu,\gamma}\left(t_\nu^{-1}(u), t_\nu^{-1}(v)\right)$	$\gamma \in (-1,1)$
Clayton	$C_\gamma^{Cl}(u,v) = \left(u^{-\gamma} + v^{-\gamma} - 1\right)^{-1/\gamma}$	$\gamma \in (0,\infty)$
Gumbel	$C_\gamma^{Gu}(u,v) = \exp\left(-\left((-\ln(u))^\gamma + (-\ln(v))^\gamma\right)^{1/\gamma}\right)$	$\gamma \in [1,\infty)$

Table 4: The definition of four popular copula families. The function Φ_γ denotes the cdf of a bivariate normal with standard margins and Pearson correlation coefficient γ and $\varPhi$ is the cdf of a standard univariate normal. t_ν denotes the cdf of a standard univariate t distribution with ν dof and $\mathbf{t}_{\nu,\gamma}$ is the cdf of a bivariate-t with standard t_ν margins and dispersion coefficient γ.

B Five Popular Dependence Measures

In Table 5, we give the definition of five dependence measures. Other dependence measures and interpretations in an actuarial context can be found in Denuit et al. (2005).

Dependence measure	Definition	Range
Spearman's rho	$\rho_S(X,Y) = 12\mathbb{E}[F(X)G(Y)] - 3$	$[-1,1]$
Kendall's tau	$\rho_\tau(X,Y) = 2\mathbb{P}[(X-\widetilde{X})(Y-\widetilde{Y}) > 0] - 1$	$[-1,1]$
Lower tail dependence	$\rho_{LT}(X,Y) = \lim_{t\downarrow 0} \mathbb{P}[F(X) \le t \mid G(Y) \le t]$	$[0,1]$
Upper tail dependence	$\rho_{UT}(X,Y) = \lim_{t\uparrow 1} \mathbb{P}[F(X) > t \mid G(Y) > t]$	$[0,1]$
Quantile exceedance probability ($\alpha \in (0,1)$)	$\rho_{Q,\alpha}(X,Y) = \mathbb{P}[F(X) > \alpha \mid G(Y) > \alpha]$	$\left[\max\left\{0, \frac{1-2\alpha}{1-\alpha}\right\}, 1\right]$

Table 5: The definition of five widely used dependence measures. $(\widetilde{X},\widetilde{Y})$ denotes an independent copy of (X,Y). Note that the limits ρ_{LT} and ρ_{UT} not necessarily exist. Note that for the shown dependence measures, we have $\rho(X,Y) = \rho(U,V)$.

C Link Functions

In Table 6, we give details on the link functions g for the copulas in Appendix A and dependence measures in Appendix B. Some link functions are not available. For instance, the link function for a Gaussian copula with upper tail dependence does not exist because a Gaussian copula C_γ^{Ga} is tail independent for $\gamma \in (-1, 1)$.

	Gaussian	t	Clayton	Gumbel
Spearman's rho	$2\sin\left(\frac{\theta\pi}{6}\right)$	NUM	NUM	NUM
Kendall's tau	$\sin\left(\frac{\theta\pi}{2}\right)$	$\sin\left(\frac{\theta\pi}{2}\right)$	$\frac{2\theta}{1-\theta}$	$\frac{1}{1-\theta}$
Lower tail dependence	NA	$\frac{2\nu+2}{\nu+1+\left(t_{\nu+1}^{-1}\left(\frac{\theta}{2}\right)\right)^2} - 1$	$-\frac{\ln(2)}{\ln(\theta)}$	NA
Upper tail dependence	NA	$\frac{2\nu+2}{\nu+1+\left(t_{\nu+1}^{-1}\left(\frac{\theta}{2}\right)\right)^2} - 1$	NA	$\frac{\ln(2-\theta)}{\ln(2)}$
Quantile exceedance probability	NUM	NUM	NUM	$\frac{\ln(2)}{\ln\left(\frac{\ln((1-\alpha)\theta-1+2\alpha)}{\ln(\alpha)}\right)}$

Table 6: Link functions as defined in Section 2 for the Gaussian, t, Clayton and Gumbel copula families combined with the dependence measures ρ_S, ρ_τ, ρ_{LT}, ρ_{UT} and $\rho_{Q,\alpha}$. Some are accessible only numerically (NUM) and some are not available (NA).

D Definition and Properties of the Beta distribution

A random variable $R \in (0, 1)$ is *Beta distributed*, if its density is given by $f_R(x) = x^{\alpha-1}(1-x)^{\beta-1} / \mathrm{B}(\alpha, \beta)$ for $x \in (0, 1)$, $\alpha, \beta > 0$ and $\mathrm{B}(\cdot, \cdot)$ denoting the Beta function. Mean and variance are given by $\mathbb{E}[R] = \alpha/(\alpha + \beta)$ and $\mathrm{var}(R) = \alpha\beta/((\alpha+\beta)^2(\alpha+\beta+1))$. The parameters can be inferred from the moments through $\alpha = \mathbb{E}[R]^2(1-\mathbb{E}[R]) / \mathrm{var}(R) - \mathbb{E}[R]$ and $\beta = \alpha(\mathbb{E}[R]^{-1} - 1)$. For $\alpha, \beta \geq 1$ the distribution is unimodal. Four Beta densities are illustrated in Figure 1.

The random variable $a + (b-a)R \in [a, b]$ for some $a < b$ is said to have a *shifted Beta distribution* with endpoints a and b.

Note that a (shifted) Beta distributed random variable R with endpoints a and b obeys $\mathrm{var}(R) < (\mathbb{E}[R] - a)(b - \mathbb{E}[R])$. Hence, for θ very close a or b, it is not possible to satisfy both $\mathbb{E}[\varphi_k|\theta] = \theta$ and $\mathrm{var}(\varphi_k|\theta) = \sigma_k^2$ with a Beta distribution.

To circumvent this problem, we propose to set $\mathrm{var}(\varphi_m|\theta) = \min\left\{\widehat{\nu_m^2}, 0.9\cdot(\theta-a)(b-\theta)\right\}$ instead of $\mathrm{var}(\varphi_k|\theta) = \sigma_k^2$.

References

Beaudoin, D., Genest, C. and Remillard, B. (2009) Goodness-of-fit tests for copulas: A review and a power study. *Insurance: Mathematics and Economics* **44(2)**, 199–213.

Berger, J. and Bernardo, J. (1992) On the development of reference priors. *Bayesian statistics* **4**, 35–60.

Bernardo, J. and Smith, A. (1994) *Bayesian Theory*. Wiley, Chichester.

BPV (2006) *Technical document on the Swiss Solvency Test*. Swiss Federal Office of Private Insurance.

Bühlmann, H. and Gisler, A. (2005) *A Course in Credibility Theory and its Applications*. Springer, Berlin.

CEIOPS (2008) *Report on its fourth quantitative impact study (QIS4) for Solvency II*. Committee of European Insurance and Occupational Pensions Supervisors.

CEIOPS (2010) *Technical specifications of its fifth quantitative impact study (QIS5) for Solvency II*. Technical report, Committee of European Insurance and Occupational Pensions Supervisors.

Clemen, R., Fischer, G. and Winkler, R. (2000) Assessing dependence: Some experimental results. *Management Science* **46(8)**, 1100–1115.

Clemen, R. and Winkler, R. (1999) Combining probability distributions from experts in risk analysis. *Risk Analysis* **19(2)**, 187–203.

Cooke, R. (1991) *Experts in Uncertainty: Opinion and Subjective Probability in Science*. Oxford University Press, New York.

Cooke, R. and Goossens, L. (2000) *Procedures guide for structured expert judgment*. Technical report, University of Technology Delft.

Cooke, R., Mendel, M. and Thijs, W. (1988) Calibration and information in expert resolution. *Automatica* **24(1)**, 87–94.

Dacorogna, M. and Canestraro, D. (2010) The influence of risk measures and tail dependencies on capital allocation. *SCOR Papers* **7**, 1–16.

Daneshkhah, A. (2004) *Uncertainty in probabilistic risk assessment: A review*. Working paper, The University Of Sheffield.

Delbecq, A., Van de Ven, A. and Gustafson, D. (1975) *Group Techniques for Program Planning: A Guide to Nominal Group and Delphi Processes*. Scott, Foresman and Company, New Jersey.

Denuit, M., Dhaene, J., Goovaerts, M. and Kaas, R. (2005) *Actuarial Theory for Dependent Risks: Measures, Orders and Models*. Wiley, Chichester.

Donnelly, C. and Embrechts, P. (2010) The devil is in the tails: Actuarial mathematics and the subprime mortgage crisis. *ASTIN Bulletin* **10(1)**, 1–33.

Eddy, D. (1982) Probabilistic reasoning in clinical medicine: Problems and opportunities. In *Judgment under uncertainty: Heuristics and biases*. (ed. Kahneman, D., Slovic, P. and Tversky, A.), pp. 249–267, Cambridge University Press, Cambridge.

Embrechts, P., McNeil, A. and Straumann, D. (2002) Correlation and dependence in risk management: Properties and pitfalls. In *Risk Management: Value at Risk and Beyond*. (ed. Dempster, M.), pp. 176–223, Cambridge University Press, Cambridge.

Embrechts, P., Mikosch, T. and Klüppelberg, C. (2001) *Modelling Extremal Events for Insurance and Finance.* Springer, Berlin.

FINMA (2008) *Swiss Solvency Test (SST), Circular 2008/44.* Eidgenössische Finanzmarktaufsicht FINMA.

Garthwaite, P., Kadane, J. and O'Hagan, A. (2005) Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association* **100(470)**, 680–701.

Genest, C., Gendron, M. and Bourdeau-Brien, M. (2009) The advent of copulas in finance. *The European Journal of Finance* **15(7–8)**, 609–618.

Genest, C. and Zidek, J. (1986) Combining probability distributions: A critique and an annotated bibliography. *Statistical Science* **1(1)**, 114–135.

Gokhale, D. and Press, S. (1982) Assessment of a prior distribution for the correlation coefficient in a bivariate normal distribution. *Journal of the Royal Statistical Society* **145**, 237–249.

Hellwig, M. (2009) Systemic risk in the financial sector: An analysis of the subprime-mortgage financial crisis. *De Economist* **157(2)**, 129–207.

Huard, D., Evin, G. and Favre, A. (2006) Bayesian copula selection. *Computational Statistics and Data Analysis* **51(2)**, 809–822.

Jouini, M. and Clemen, R. (1996) Copula models for aggregating expert opinions. *Operations Research* **44(3)**, 444–457.

Kahneman, D. and Tversky, A. (1973) On the psychology of prediction. *Psychological review* **80(4)**, 237–251.

Kahneman, D. and Tversky, A. (1982) On the study of statistical intuitions. *Cognition* **11**, 123–141.

Kallen, M. and Cooke, R. (2002) Expert aggregation with dependence. In *Probabilistic Safety Assessment and Management*, pp. 1287–1294, International Association for Probabilistic Safety Assessment and Management.

Kynn, M. (2005) *Eliciting expert knowledge for bayesian logistic regression in species habitat modelling*. PhD thesis, School of Mathematical Sciences, Queensland University of Technology.

Kynn, M. (2008) The "heuristics and biases" bias in expert elicitation. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **171(1)**, 239–264.

Lambrigger, D., Shevchenko, P. and Wüthrich, M. (2007) The quantification of operational risk using internal data, relevant external data and expert opinions. *Journal of Operational Risk* **2(3)**, 3–27.

Lichtenstein, S. and Fischhoff, B. (1980) Training for calibration. *Organizational Behavior and Human Performance* **26**, 149–171.

Lindley, D. (1965) *Introduction to Probability and Statistics from a Bayesian Viewpoint*. Cambridge University Press, Cambridge.

Linstone, H. and Turoff, M. (1975) *Delphi Method: Techniques and Applications*. Addison-Wesley, London.

McNeil, A., Frey, R. and Embrechts, P. (2005) *Quantitative Risk Management: Concepts, Techniques, Tools*. Princeton University Press, Princeton.

Meyer, M. and Booker, J. (2001) *Eliciting and Analyzing Expert Judgment: A Practical Guide*. Society for Industrial and Applied Mathematics, Philadelphia.

Morgan, M. and Henrion, M. (1992) *Uncertainty: A Guide to Dealing With Uncertainty in Quantitative Risk and Policy Analysis.* Cambridge University Press, Cambridge.

Mosleh, A., Bier, V. and Apostolakis, G. (1988) A critique of current practice for the use of expert opinions in probabilistic risk assessment. *Reliability Engineering and system safety* **20(1)**, 63–85.

Nelsen, R. (2006) *An Introduction to Copulas. 2nd edition*. Springer, New York.

O'Hagan, A., Buck, C., Daneshkhah, A., Eiser, J., Garthwaite, P., Jenkinson, D., Oakley, J. and Rakow, T. (2006) *Uncertain Judgements: Eliciting Experts' Probabilities*. Wiley, Chichester.

Ouchi, F. (2004) *A literature review on the use of expert opinion in probabilistic risk analysis*. Working Paper 3201, World Bank Policy Research.

Price, H. and Manson, A. (2002) Uninformative priors for Bayes' theorem. *Bayesian Inference and Maximum Entropy Methods in Science and Engineering* **617(1)**, 379–391.

Rebonato, R. (2010) *Coherent Stress Testing*. Princeton University Press, Chichester.

Robert, C. and Casella, G. (2005) *Monte Carlo Statistical Methods*. Springer, New York.

Schnieper, R. (2010) The Use of Models for Solvency Testing Purposes. Presentation at the "Fest-Symposium 'Versicherungsmathematik im Wandel' Zum 80. Geburtstag von Prof. Dr. Hans Bühlmann".

Tversky, A. and Kahneman, D. (1974) Judgment under uncertainty: Heuristics and biases. *Science* **185(4157)**, 1124–1131.

Van der Vaart, A. (1998) *Asymptotic Statistics*. Cambridge University Press, Cambridge.

Wallsten, J. and Budescu, D. (1983) Encoding subjective probabilities: A psychological and psychometric review. *Management Science* **29**, 151–173.

Winkler, R. (1968) The consensus of subjective probability distributions. *Management Science* **15**, B61–B75.

SCOR Papers

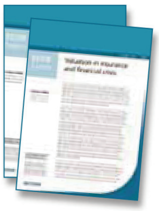

SCOR Paper N°1 - September 2008
Using Capital Allocation to Steer the Portfolio towards Profitability

SCOR Paper N°2 - Août 2008
La bancassurance : généralisation ou déclin du modèle ?

SCOR Paper N°3 - December 2008
Valuation in insurance and financial crisis

SCOR Paper N°4 - March 2009
Modern Companies and Extreme Risks

SCOR Paper N°5 - July 2009
Securitization, Insurance and Reinsurance

SCOR Paper N°6 - January 2010
Adapting the solvency regulations to times of crisis, accepting the riskiness of the situation

SCOR Paper N°7 - January 2010
The Influence of Risk Measures and Tail Dependencies on Capital Allocation

SCOR Paper N°8 - March 2010
Principle-based Solvency: A Comparison between Solvency II and the Swiss Solvency Test

SCOR Paper N°9 - December 2010
Aging: a Global Phenomenon

SCOR Paper N°10 - December 2010
PrObEx: a New Method for the Calibration of Copula Parameters from Prior Information, Observations and Expert Opinions

SCOR Paper N°11 - December 2010
Development Risk

SCOR Paper N°12 - Décembre 2010
Study of the impact of inflation and GDP growth on property-liability and life insurance premiums over the last 30 years: case of the G7 countries

Global Risk Center

SCOR Papers, edited by SCOR, are one of the tool supporting the SCOR Global Risk Center.
The SCOR Global Risk Center gathers and analyses the most interesting resources about risks. It operates as a dual resource center, based both on data and resources produced by SCOR itself, and on all other resources available selected specifically by SCOR. Publications will be available in English, French and/or German.
SCOR Global Risk Center is available at www.scorglobalriskcenter.com or on SCOR's website – www.scor.com.