

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Sequential Refined Partitioning for probabilistic dependence
assessment

May 14, 2018

For Peer Review

Abstract

Modelling dependence probabilistically is crucial for many applications in risk assessment and decision making under uncertainty. Neglecting dependence between multivariate uncertainties can distort model output and prevent a proper understanding of the overall risk. Whenever relevant data for quantifying and modelling dependence between uncertain variables is lacking, expert judgement might be sought to assess a joint distribution. Key challenges for the use of expert judgement for dependence modelling are over- and underspecification. An expert can sometimes provide assessments which are not consistent with any probability distribution (overspecification), and on the other hand, without making very restrictive parametric assumptions an expert cannot fully define a full probability distribution (underspecification). The Sequential Refined Partitioning method addresses over- and underspecification whilst allowing for flexibility about which part of a joint distribution is assessed and its level of detail. Potential overspecification is avoided by ensuring low cognitive complexity for experts through eliciting single conditioning sets and by offering feasible assessment ranges. The feasible range of any (sequential) assessment can be derived by solving a linear programming problem. Underspecification is addressed by modelling the density of directly and indirectly assessed distribution parts as minimally informative given their constraints. Hence, our method allows for modelling the whole distribution feasibly and in accordance with experts' information. A non-parametric way of assessing and modelling dependence flexibly in such detail has not been presented in the expert judgement literature for probabilistic dependence models so far. We provide an example of assessing terrorism risk in insurance underwriting.

Keywords: Structured Expert Judgement, Dependence Modelling, Minimum Information, Terrorism Risk, Uncertainty Modelling, Risk Analysis

1 INTRODUCTION

1 In many risk and decision analysis problems, we need to quantify uncertainties and their dependence
2 as otherwise a model for risk assessment and decision making might not be fit for purpose. Indeed,
3 quantifying dependence for probabilistic modelling is listed repeatedly among the most significant
4 topics which decision and risk analysis research faces [1, 2]. Therefore, modelling joint distributions
5 in various ways and for several problem types is an active research area (e.g. Durante and Sempi
6 [3], Hanea et al. [4], McNeil et al. [5], Joe [6], Genest et al. [7], Kurowicka and Cooke [8], Embrechts
7 et al. [9]). A common challenge is a lack of relevant data for quantifying dependence models. In such
8 cases, this information should be assessed through expert judgements. A *structured* expert judge-
9 ment (SEJ) elicitation is the most sensible solution to missing historical data whenever a simplifying
10 assumption, such as independence, is not applicable. Werner et al. [10] and Werner et al. [11] discuss
11 expert judgement methods for dependence in more detail. The former outlines how it is used for
12 several dependence models and reviews commonly elicited forms together with their implication on
13 experts' cognitive burden. The latter presents the main steps of structured dependence elicitations
14 and reviews the most prevalent cognitive fallacies for assessing dependence as their mitigation is a
15 main aim of structured processes. **Most applications discussed in these reviews are based on Cooke
16 and Goossens [12, 13] which are among the first guides on SEJ procedures for dependence. Both
17 guides are of further relevance for this paper as they consider in particular the elicitation of condi-
18 tional exceedance probabilities, an elicited form we will address in more detail later.** In this paper,
19 we focus mainly on the process for quantitative elicitation, though we do discuss an approach to
20 structuring experts' knowledge prior to elicitation in the illustrative example of section 4.

21 For us, dependence means that multiple uncertainties are present and obtaining information about
22 one changes the uncertainty assessment of the other(s). More specifically, we consider the bivariate
23 dependence between two random variables X and Y with joint distribution function $F_{X,Y}(x, y)$ and
24 marginal distributions $F_X(x)$ and $F_Y(y)$. The variables are independent if the assessment of Y does
25 not change when given information about X . Dependence is simply the absence of independence. It
26 is a property of experts' knowledge (and beliefs) and its definition falls therefore into the subjective
27 probability context as in line with De Finetti [14], Savage [15] and Ramsey [16].

28 We address the problem that experts can only ever assess certain aspects of a joint distribution
29 whereas a decision-maker might desire these assessments to be made at a detailed level. The former
30 implies that we have a partially unknown distribution for which various alternatives fit the given
31 information. This is known as model *underspecification*. More specifically, we are only ever given the
32 probability mass (or density) within some distribution parts, either through their direct assessment
33 or (in parts which are never assessed) through the indirect result of these parts together with re-

lated assessed parts having to comply with the marginals. However, we can model these probability masses in various forms which all have the right amount (i.e. are feasible). Of course, we might elicit additional information from experts to distinguish between distributions, yet we need to acknowledge the impossibility of ever eliciting enough information to single out a unique distribution. This is unless adopting a low-dimensional parametric model early on in the modelling process¹. Such parametric assumptions nevertheless restrict the obtained knowledge on dependencies and we might miss potentially important model aspects, such as random variables' behaviour in the extreme parts (tails) of a joint distribution. Hence, it is often desirable to avoid distributional assumptions which might exclude phenomena that the expert thinks are important.

Within a non-parametric setting, an elicitation should capture detailed distribution features, e.g. the probability mass within narrowly defined parts of the distribution, such as the tails to determine tail dependence, as they result in a more specific distribution, thus making the model more valuable for a decision-maker. Nevertheless, while detailed assessments might be desired by decision-makers, they increase the experts' cognitive burden, potentially resulting in inconsistent and infeasible assessments. This is termed *overspecification*, the second modelling challenge that we encounter².

As a non-parametric approach, addressing under- and overspecification, we present the *sequential refined partitioning* (SRP) method for assessments that can be made to any level of detail for any part of a joint distribution. In the SRP method, we address overspecification through an elicitation procedure which never increases the conditioning set to more than one condition and thus maintains a low cognitive complexity. Further, the procedure ensures consistent and feasible assessments through explicit guidance on assessments' feasibility ranges. Underspecification is dealt with by allowing the expert to specify as much detail as is desired and by then determining the density form of directly and indirectly assessed parts of the distribution through the unique copula distribution that is minimally informative with respect to the independent copula and that corresponds to the elicited information. This makes only the weak assumption that in the absence of any specific guidance from the expert we should make the copula as close as possible to the independent copula (in the sense of minimizing information). This ensures that the whole distribution is in agreement with the experts beliefs.

We note that there may be other situations where a joint distribution is to be defined but data is incomplete. For example there may be few data, and/or there may be few or no data in the tails. In these cases, we can apply SRP as part of a hybrid method for dependence assessment, so that it can also be applied for copula model selection more generally, i.e. in the research area of empirical

¹Under low-dimensional parametric assumptions, it suffices to assess a chosen form's main parameters. E.g. eliciting the mean vector and the covariance matrix quantifies a multivariate Gaussian distribution sufficiently.

²Overspecification can also occur with parametric models, e.g. if assessed covariances jointly do not result in a positive definite matrix.

66 **copula estimation.**

67 The minimum information approach offers a recognised approach to incomplete knowledge[17]. Fur-
68 ther, it allows us to stop the elicitation process at any time and still derive a unique distribution
69 (in contrast to common probabilistic dependence models for which a full conditional probability
70 table is required, e.g. Bayesian (Belief) nets (BNs) [18]). In the context of dependence elicita-
71 tion, minimum information methods (and related approaches) have been used before, for instance
72 in probabilistic inversion (PI) methods [19, 20, 21, 22], vine-copula quantification [23, 24], or as
73 well joint distributions more generally within decision analysis contexts [25, 26, 27, 28]. However,
74 these previous methods do not consider flexible nor detailed (e.g. tail) dependence assessments and
75 their impact on potential overspecification of experts' judgements and on the minimum information
76 solution to underspecification. For example, Bedford et al. [19] explicitly provide guidance on fea-
77 sibility constraints. Yet, they consider dependence elicitation at a rather broad level, eliciting only
78 a small number of assessments. This restricts the information to be obtained already early on in
79 the modelling process and thus neglects focusing on specific parts of a distribution more exclusively.
80 The SRP method's contribution is therefore that we provide an elicitation procedure to assess any
81 part of a distribution to any desired level of detail while maintaining low cognitive complexity and
82 avoiding infeasible expert judgements. As such, it also contributes to expert judgement methods
83 for dependence in which increasing conditioning sets pose a concern (see Werner et al. [10] for a
84 discussion). Similarly, the SRP method's approach to underspecification is more detailed than in
85 previous research.

86 **These contributions emphasise the applicability of our method in higher dimensions more gener-**
87 **ally. While in this paper we focus on assessing bivariate dependence, it should be noted that any**
88 **d -dimensional copula density can be built through $d(d - 1)/2$ so-called pair (bivariate) copulas**
89 **through a vine structure [24]. This method of modelling dependence, in conjunction with appropri-**
90 **ate simplifying assumptions, can avoid the *curse of dimensionality* (see Nagler and Czado [29] for**
91 **using this approach in the context of kernel estimation). In that way, our method can be extended**
92 **to higher dimensions of dependence and be used more generally in the area of multivariate density**
93 **estimation. As such, the SRP method can contribute to more traditional methods of copula esti-**
94 **vation for tail dependence assessment through experts when data on extremes are rare.**

95 Figure 1 illustrates the method's modelling context schematically.

96 In the upper part, we observe that incomplete knowledge leads inevitably to an underspecified
97 model. This is solved by a minimum information approach. In order to derive a model that is
98 valuable for a decision-maker, the modelling process deviates along the dashed lines to the lower part.
99 Here, the constraints of the minimum information problem determined by the experts' judgements

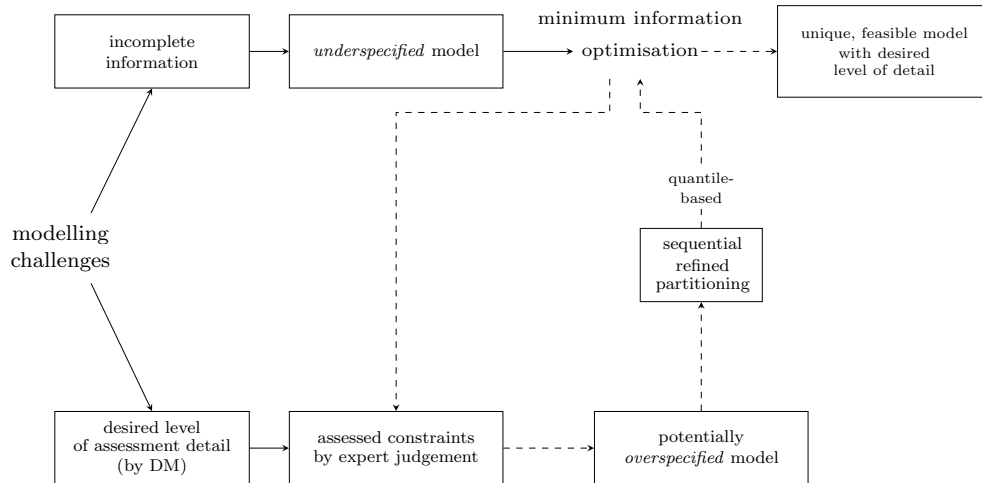


Figure 1: Modelling context of the SRP method.

are assessed as detailed as desired. As these might be overspecified, we use an elicitation process that leads to feasible assessments. In the remainder of this paper, this is presented in section 2, introducing the elicitation procedure, and section 3, outlining the optimisation problem. Section 4 shows how our method has been used in an insurance underwriting risk assessment of political violence/terrorism in which a detailed and flexible method is of particular interest for stress-testing a model. Finally, section 5 concludes the paper.

2 ELICITING DETAILED DEPENDENCE INFORMATION FEASIBLY AND CONSISTENTLY THROUGH *SEQUENTIAL REFINED PARTITIONING*

In this section, we introduce our sequential elicitation procedure which addresses the potential issue of overspecification by providing explicit guidance on making feasible and consistent assessments. In the expert judgement literature, several approaches to ensuring feasibility and consistency are proposed, each with different implications on the robustness of the final assessment result. As such, some methods (always) allow for an assessment within the elicited forms' standard ranges (for correlation coefficients $\in [-1, 1]$ and for conditional and joint probabilities $\in [0, 1]$). However, this might jeopardise experts' commitment and confidence in the elicitation method if assessments are adjusted afterwards (for ensuring feasibility). While other methods do not modify assessments, they might increase experts' cognitive complexity. For instance, by limiting assessment ranges (away from the aforementioned standard ones), or by imposing unrealistic assumptions onto experts' understanding of elicited forms, e.g. when eliciting conditional judgements with large conditioning sets. For the

1
2
3
4 117 latter, we might expect an expert to include and equally consider all the information given by a
5
6 118 large conditioning set so that common cognitive fallacies, such as the conjunction fallacy and its
7
8 119 conditional version (see Werner et al. [11] for an overview on heuristic and biases in dependence
9
10 120 assessment), should be (ideally) avoided and hence feasibility is given. Yet, this might not be guar-
11
12 121 anteed.

12 122 In our method, we do not impose such unrealistic assumptions on experts' cognitive capabilities, nor
13
14 123 do we modify assessments after they have been given. Rather, we only ever elicit single conditioning
15
16 124 sets and give guidance on possible feasible assessment ranges. This includes not only providing the
17
18 125 corresponding upper and lower bounds but also explaining their interpretation.

19 126 Mathematically, the feasibility range for any sequential assessment procedure is derived by solving
20
21 127 a linear programming (LP) problem (see Vanderbei [30] for an introduction to LP). The number
22
23 128 of constraints is restricted to a maximum of nine, irrespective of the number of elicitations. In the
24
25 129 remainder of this section, we first present the general set-up together with the relevant proofs before
26
27 130 we outline some specific elicitation sequences, which we regard as of interest for several practical
28
29 131 applications.

132 2.1 General set-up of sequentially refined partitioning

133 We shall start by introducing some definitions. The unit square is here defined as the product of
134
135 $(0, 1] \times (0, 1]$. Given values $u_0 = 0 < u_1 < \dots < u_n < 1 = u_{n+1}$, and $v_0 = 0 < v_1 < \dots < v_m < 1 =$
136
137 v_{m+1} , we define the associated *quantile partition* of the unit square as the set of rectangles of the
138
139 form $(u_i, u_{i+1}] \times (v_i, v_{i+1}]$. We call this set of rectangles $QP(u, v)$.

140 Given (p, q) with p different to the u_i and q different to the v_j , the (p, q) -refinement of $QP(u, v)$,
141
142 denoted $QP(u, v; p, q)$, is the quantile partition obtained by including p and q in the values for u
143
144 and v respectively. All rectangles in the old partition are either in the new partition or are a union
145
146 of two or four rectangles of the old partition. Figure 2 shows two partitioned example distributions
147
148 which result from any number of previously elicited quantiles (solid lines) in addition to new ones
149
150 (dashed lines).

151 A probability distribution on a quantile partition $QP(u, v)$ simply assigns a probability value
152
153 to each rectangle of the quantile partition. A (p, q) -refinement of such a probability distribution is
154
155 a probability distribution on $QP(u, v; p, q)$ such that the probability of a rectangle in $QP(u, v)$ is
156
157 either the same as it is in the (p, q) -refinement of $QP(u, v)$, or it equals the sum of the probabilities
158
159 of the rectangles that make it up.

160 A *merging* of a quantile partition $QP(u; v)$ is obtained by merging together some of the partition
161
162 rectangles in such a way that we still have a quantile partition. This can also be obtained by taking




Figure 2: Partition example $QP(\tilde{u}, \tilde{v}; p, q)$ with solid lines for previously elicited quantiles and dashed lines for new ones.

a subsequence of the u 's and v 's and building the corresponding quantile partition. A *merged probability distribution* on the refined quantile partition is obtained by adding together the probabilities of the rectangles in each refined rectangle.

We always work with *discrete copula distributions*, which are probability distributions on a quantile partition that have the additional property that (for any k) the sum of probabilities of rectangles $(u_i, u_{i+1}] \times (v_i, v_{i+1}]$ with $u_{i+1} \leq u_k$ is equal to u_k , and similarly, the sum of all probabilities of rectangles $(u_i, u_{i+1}] \times (v_i, v_{i+1}]$ with $v_{i+1} \leq v_k$ is equal to v_k . For a general introduction to copula theory, see Nelsen [31], Joe [6] and Durante and Sempi [3]. However, note that most theory is on continuous copulas with marginals being continuous uniform distributions. For an overview on elicitation methods for copulas, see Werner et al. [10].

Proposition 1. *Suppose we are given values $u_0 = 0 < u_1 < \dots < u_n < 1 = u_{n+1}$, and $v_0 = 0 < v_1 < \dots < v_m < 1 = v_{m+1}$ (where $n, m > 0$), $0 < p, q < 1$, with p different to the u_i and q different to the v_j . Then a copula distribution on $QP(u, v)$ can be refined to a copula distribution on $QP(u, v; p, q)$.*

The proof of proposition 1 is found in the Appendix.

Having shown that we can always refine a copula distribution as above, we now wish to establish the possible range of values that can be taken by the rectangle $(p, 1] \times (q, 1]$ in a refined copula distribution. That is, we depart from the specific copula refinement defined in the Proof of Proposition 1, and ask what range of values can be allocated as the probability of $(p, 1] \times (q, 1]$ in some copula refinement.

Suppose that i and j are chosen such that u_i is the largest of the u -quantiles that is smaller than p , and v_j is the largest of the v -quantiles that is smaller than q (this includes the possibility that u_i or v_j is 0, or that u_{i+1} or v_{j+1} is 1). Define $\tilde{u}_1 = u_i$, $\tilde{u}_2 = u_{i+1}$, $\tilde{v}_1 = v_j$ and $\tilde{v}_2 = v_{j+1}$. The quantile




Figure 3: Maximum case of 16 partitions (right) resulting from partitioning 9 rectangles (left).

174 partition $QP(\tilde{u}, \tilde{v})$ is a merging of $QP(u, v)$, and we can merge the copula distribution on $QP(u, v)$
 175 to get one on $QP(\tilde{u}, \tilde{v})$.

176 Furthermore $QP(\tilde{u}, \tilde{v}; p, q)$ is a merging of $QP(u, v; p, q)$. Note that $QP(\tilde{u}, \tilde{v})$ has at most 9 rectan-
 177 gles and that $QP(\tilde{u}, \tilde{v}; p, q)$ has at most 16 rectangles - see Figure 3.

178 For convenience we shall now consider only the case of 16 rectangles, which occurs when $u_i, v_j \neq 0$
 179 and $u_{i+1}, v_{j+1} \neq 1$, as shown on the right of Figure 3. Other cases are simplifications of the one we
 180 consider here and can be dealt with in the same way.

181 We label the 16 rectangles of $QP(\tilde{u}, \tilde{v}; p, q)$ as R_{11}, \dots, R_{44} as shown in the right hand of Figure 3.

The 9 rectangles of $QP(\tilde{u}, \tilde{v})$ are labelled as $\tilde{R}_{11}, \dots, \tilde{R}_{3,3}$ as shown in the left hand of Figure 3.
 Clearly $R_{11}, \dots, R_{4,4}$ are each unions of rectangles in $QP(u, v)$, and furthermore,

$$R_{12} \cup R_{13} = \tilde{R}_{12}$$

$$R_{42} \cup R_{43} = \tilde{R}_{32}$$

$$R_{21} \cup R_{31} = \tilde{R}_{21}$$

$$R_{24} \cup R_{34} = \tilde{R}_{23}$$

$$R_{22} \cup R_{23} \cup R_{33} \cup R_{32} = \tilde{R}_{22}.$$

182 Suppose we are given a copula distribution on $QP(\tilde{u}, \tilde{v})$, for which \tilde{p}_{st} is the probability of \tilde{R}_{st}
 183 ($s, t = 1, 2, 3$). We wish to assign copula probabilities p_{st} to the rectangles R_{st} ($s, t = 1, 2, 3, 4$) so
 184 that the new distribution merges to p on $QP(\tilde{u}, \tilde{v})$.

185 For the merging we simply require,

- 186 • for the corner rectangles of $QP(\tilde{u}, \tilde{v})$: $p_{11} = \tilde{p}_{11}, p_{14} = \tilde{p}_{13}, p_{41} = \tilde{p}_{31}, p_{44} = \tilde{p}_{33}$,

- 1
2
3
4 187 • for the central rectangle in $QP(\tilde{u}, \tilde{v})$: $p_{22} + p_{32} + p_{23} + p_{33} = \tilde{p}_{22}$,
5
6 • for the remaining rectangles
7
8

$$9 \quad p_{12} + p_{13} = \tilde{p}_{12}$$

$$10 \quad p_{42} + p_{43} = \tilde{p}_{32}$$

$$11 \quad p_{21} + p_{31} = \tilde{p}_{21}$$

$$12 \quad p_{24} + p_{34} = \tilde{p}_{23}.$$

13
14
15
16
17
18 To ensure that the new distribution is a copula we also need to impose two constraints corresponding
19 to a row and a column:
20

$$21 \quad p_{21} + p_{22} + p_{23} + p_{24} = p - \tilde{u}_1$$

$$22 \quad p_{12} + p_{22} + p_{32} + p_{42} = q - \tilde{v}_1.$$

23
24
25
26
27 (Note that these constraints correspond to row 2 and column 2 of the right hand of Figure 3. We
28 could also have specified similar constraints on row 3 and column 3, but it straightforward to see
29 that these are redundant).
30
31

32 Now define,
33

$$34 \quad f(p_{11}, \dots, p_{44}) = p_{33} + p_{43} + p_{34} + p_{44}$$

35
36
37 188 to be the total probability in the square $(p, 1] \times (q, 1]$. This is a linear function of the p_{st} and we are
38 free to choose it to take any value subject to the constraints listed above. As all these are linear,
39 we immediately see that we have the form of a linear programming problem, and so the range of
40 allowable values is an interval whose maximum and minimum values can be found by solving 2 LP
41 problems. The cases in which $QP(\tilde{u}, \tilde{v}; p, q)$ has fewer than 16 rectangles work similarly. The above
42 discussion (with minor adaptations to the other cases by removing further redundant constraints)
43 can be summarized in the following Proposition:
44
45
46
47 194

48
49 **Proposition 2.** *The range of feasible values for the probability of $(p, 1] \times (q, 1]$ in any copula
50 refinement of the copula distribution on $QP(\tilde{u}, \tilde{v})$ is given by the interval:*
51

$$52 \quad [\min f, \max f],$$

53
54
55
56
57 195 given the corresponding constraint sets.

58 196 We can obtain $\min f$ and $\max f$ by solving feasible LP problems with at most 12 variables and 9
59 constraints.
60 197

This now allows us to construct an algorithm for assessing copulas with expert judgements for quantile exceedance probabilities of the form:

$$P(Y > y_q | X > x_p)$$

where x_p and y_q index the p^{th} and q^{th} quantile for X and Y accordingly. For example, $p = 0.5$ and $q = 0.5$ correspond to the medians of X and Y . Other distribution areas can then be derived. Given a number of such coherent elicitation at quantile pairs $(u_1, v_1), \dots, (u_n, v_n)$ we can calculate the copula distribution on the copula partition $QP(u, v)$.

For a new quantile pair (p, q) , we then solve the LP problem to obtain the exact feasible range for the probability of $(p, 1] \times (q, 1]$. Note that this does not fully specify the distribution on all elements of the refined partition $QP(u, v; p, q)$. To achieve this, either

- (a) we can carry out further elicitation at corner points in $QP(u, v; p, q)$ using proposition 2 repeatedly for obtaining feasible ranges from the expert; or
- (b) we can make assumptions, such as minimally informative probabilities to restrict the number of elicitation required.

In the next section, we give a simple example of making assessments in the tail of the distribution along the lines of (a) but carried out in a slightly different order as there are few constraints in this case.

2.2 Commonly assessed quantile partition sequences

After having presented the mathematical set-up of refined partitioning generally, we now discuss some partitions that might be commonly assessed in practice.

One recurrent way of refining a joint distribution's assessments is by sequentially choosing a quantile for p and/or q that is always either higher or lower than any previously assessed value. Then, we elicit the corresponding area above a previously elicited quantile for a new maximum or below it for a new minimum. Such sequences assess in particular the distribution tails more explicitly.

Alternatively, it is (also) possible in our method to elicit probabilities of specific values, e.g. for 1, 10, \dots , 100 (units of elicited variable) rather than common quantiles, such as the median, if this can increase intuitiveness in particular for more extreme parts in the distribution tails. This relates to the choice of whether to frame the elicitation question in terms of quantiles or values. Both have been suggested (as P- and V-methods) since the pioneering probability elicitation by the Stanford Research Institute in the 1970s [32]. While a more recent discussion on this choice is given in [33], we consider in the following the elicitation of quantiles.

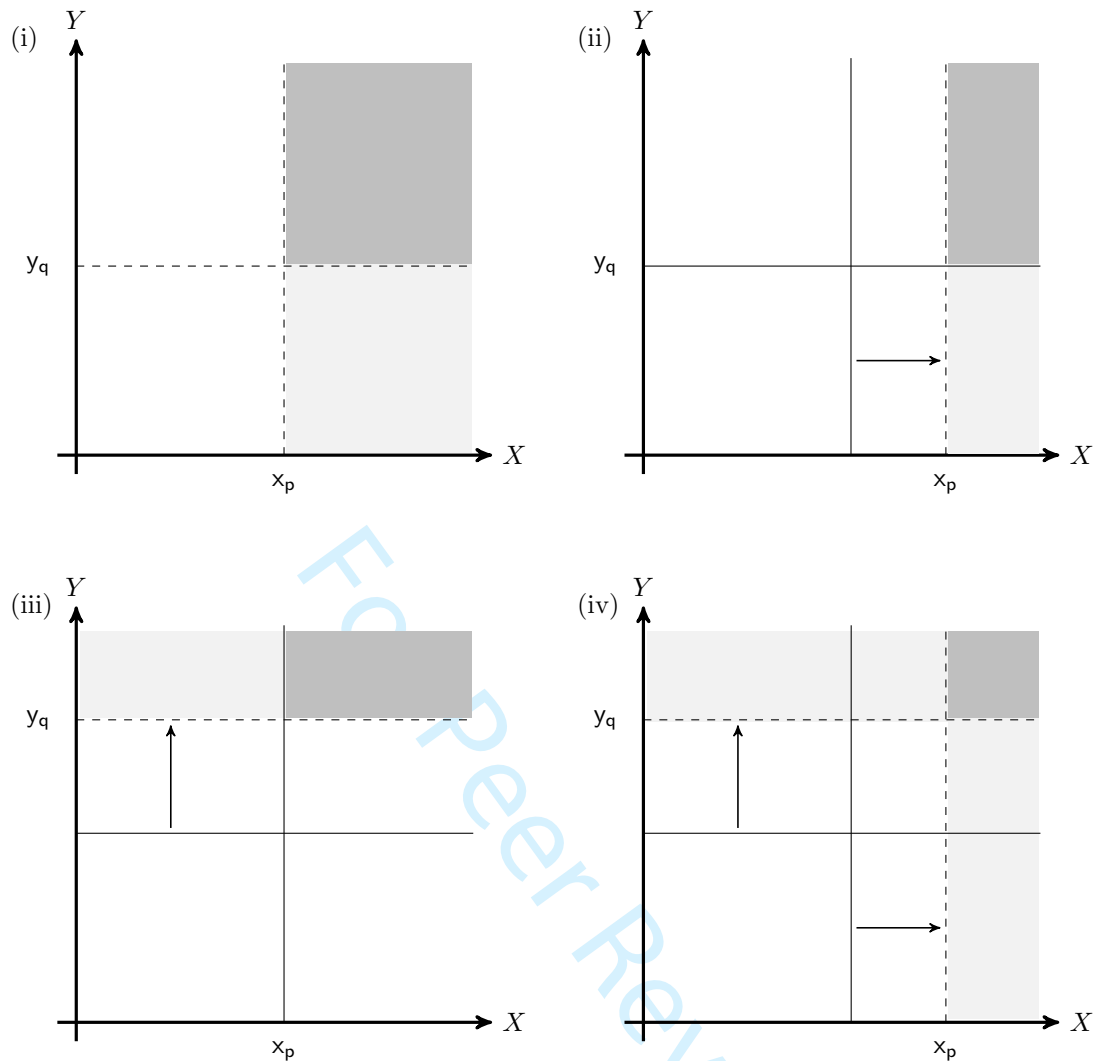


Figure 4: Example of a quantile partition for assessing the upper tail.

Figure 4 illustrates a sequence of quantile partitions on the upper tail constructed through setting new quantile maxima in (ii) to (iv) following an initial assessment (i) (note that this carries out the option (a) described in the previous section). We consider the procedure of Figure 4, i.e. further partitioning that probability mass which has been assessed directly in step (i) as most intuitive and practically useful. Nevertheless, the initial assessment also determines the probability mass in areas of the joint distribution which are not assessed further, $P(Y > y_q | X \leq x_p)$, $P(Y \leq y_q | X > x_p)$ and $P(Y \leq y_q | X \leq x_p)$, meaning we can also use a similar procedure to refine these.

First (in (i)), we elicit an overall probability mass and then subsequently refine the assessment. Suppose we first elicit $P(Y > y_{0.5} | X > x_{0.5})$. Following (i), we elicit a refined quantile partition as determined by a new x_p in (ii). A common choice here might be the 90th or 95th quantile in order to assess the probability mass in the joint




Figure 5: Quantile partition of the joint distribution from (i) to (iv).

distribution's extreme (tail) region. Thus, we elicit for instance $P(Y > y_{0.5}|X > x_{0.95})$. In the illustrative case-study of Section 4, we use a scenario mapping method [34] prior to the elicitations in order to gauge experts' familiarity with such tail judgements and decide on a quantile for which experts are comfortable to make assessments.

In (iii), we condition on Y and the new y_q is chosen to assess the tail region. With x_p being the median, we thus elicit $P(X > x_{0.5}|Y > y_{0.95})$. Depending on the underlying meaning of the variables and knowledge about causal or probabilistic relationships (see e.g. Rottman and Hastie [35], Werner et al. [11]), the expert might find it easier to condition on one variable rather than the other. Our method is flexible enough to allow for this.

In the last step of this quantile partition sequence, experts assess either $P(Y > y_{0.95}|X > x_{0.95})$ or $P(X > x_{0.95}|Y > y_{0.95})$, depending on case-specific interest, whereas p and q are the ones from the previous two rounds. Thus we further explore the joint tail region. Figure 5 displays the refinement in the quantile partition from the first to the latest assessment.

The assessments' feasibility ranges are as follows. The assessment in (i) is unrestricted, meaning experts can assess any value between $[0, 1]$. If the expert believes the variables are independent, the assessment is equal to $P(Y > y_q)$, that is learning about X does not change experts' belief. For negative dependence, the assessment is between $[0, P(Y > y_q))$ and for positive dependence, it is within $(P(Y > y_q), 1]$.

All following assessments on the other hand are restricted and only feasible if the assessed value falls within the range which is determined by solving the LP problem of minimising and maximising the possible values of the assessed area subject to the constraints that any new partition simply adds up to their previous assessments (see medium and dark grey areas \tilde{P}_k in Figure 2) while areas which

have not been newly partitioned do not change (see light grey areas \tilde{P}_k in Figure 2). Consider for example the assessment in (iv). It is only feasible within the range that is determined by solving the following LP problem (with regards to Figure 4 on the right):

$$\left. \begin{array}{l} \min \\ \max \end{array} \right\} p_{33} \quad (2.1.1)$$

subject to

$$p_{13} + p_{12} = \tilde{p}_{12} \quad (2.1.2)$$

$$p_{23} + p_{22} + p_{33} + p_{32} = \tilde{p}_{22} \quad (2.1.3)$$

$$p_{11} = \tilde{p}_{11} \quad (2.1.4)$$

and

$$p_{21} + p_{31} = \tilde{p}_{21} \quad (2.1.5)$$

Experts express negative dependence again through a judgement close or equal to the lower bound, positive dependence is expressed by judgements close or equal to the upper bound and independence is assessed as before. As the upper and/or lower bounds deviate from the standard range of $[0, 1]$, it is necessary to communicate these restricted feasibility bounds to an expert and explain their interpretation.

The procedure for assessments (ii) to (iv) is repeated as often as necessary (with appropriate modifications) to obtain a desired level of detail (see assessments (v) to (vii) in Figure 6 for the next round of three assessments). Having assessed previously the 90th or the 95th quantile of X and Y , we now might consider the 99th quantile. This allows for "zooming in" on the joint distribution's tail even further.

The resulting quantile partitions are illustrated in Figure 7.

While this section presents an example with a focus on refining the upper distribution tail, remember that the generality of the method (as introduced in Section 2.1) allows for any further refinement of the distribution, such as for instance shown in Figure 2 (on the right).

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60




Figure 6: Further refining the assessment on the joint upper distribution tail.




Figure 7: Resulting quantile partitions after further refining the previous assessments.

3 MODELLING THE FORM OF DIRECTLY AND INDIRECTLY ASSESSED PROBABILITY MASSES THROUGH MINIMUM INFORMATION

After having presented the elicitation procedure, which allows for feasibly assessing the probability mass within any part of the joint distribution, in this section we outline how we model the form of directly and indirectly assessed parts as minimally informative.

The reason for a minimum information approach is to address the modelling issue of underspecification. We do not have enough information for choosing a distribution that fits the experts' assessments uniquely but we wish to find the simplest distribution that matches them. This approach allows us to derive a unique distribution regardless the quantile partition's level of detail. As such, it does not restrict the flexibility of the assessment procedure from section 2.

Formally, we aim for modelling dependence through that copula which is chosen to have minimum information (also called *Kullback-Leibler divergence* [36]) with respect to the uniform copula given the quantile constraints. The resulting distribution is considered the most independent copula satisfying the constraints.

Consider the joint distribution $g(x, y)$ with marginal densities $g_1(x)$ and $g_2(y)$. Whenever g_1 and g_2 are not independent, i.e. $g(x, y) \neq g_1(x)g_2(y)$, we need to model the dependence between them. To do so, we introduce the concept of relative information $I(g; h)$ which is a measure of similarity between the two distributions and it is defined for $g(x)$ with respect to $h(x)$ as:

$$I(g; h) = \int g(x) \log \left(\frac{g(x)}{h(x)} \right) dx$$

Whenever $g(x) = h(x)$, it follows that $I(g; h) = 0$. A higher value of $I(g_1; g_2)$ corresponds to less similarity. We consider $h(x)$ a background distribution, commonly chosen as uniform or log-uniform. Alternatively, we use sensitivity analysis for selecting an appropriate form [19]. Together with the constraints, this choice determines the form of $g(x)$ in absence of further information [23].

Information is invariant under monotone transformations. Therefore, if c_g and c_h are copula densities associated with the previous densities g and h , we have $I(c_g; c_h) = I(g; h)$. In particular if h is the joint independent distribution with the same marginals as g (g_1 and g_2), so that $h = g_1g_2$ then $I(g; g_1g_2) = I(c; \mathbf{uniform})$ where h is the uniform copula. This gives the interpretation of our minimum information copula as the most independent copula given the constraints.

See Bedford and Wilson [37] for a detailed derivation on how a minimum information distribution

1
2
3
4 292 can be approximated by the equivalent distribution of maximum entropy³ [39].
5
6 293 For an extensive discussion on obtaining a minimum information copula through the convex optimi-
7
8 294 sation problem, we refer to [23, 37, 19]. Here, it suffices to say that the conditional density within
9
10 295 each rectangle is uniform. As discussed in Section 2, when we stop eliciting information from ex-
11
12 296 perts, some rectangles' density has been directly assessed by an expert while for other rectangles the
13
14 297 mass is given indirectly through related assessment and the marginals. In order to obtain a unique
15
16 298 solution for the whole distribution, we hence need to solve the minimisation problem of equation
17
18 299 for directly and indirectly assessed parts.
19
20 300 We refer to Bedford and Wilson [37] and Meeuwissen and Bedford [21] for the corresponding proofs
21
22 301 that such a minimum information distribution exists and is unique. Furthermore, Bedford et al. [19]
23
24 302 and Bickel and Smith [28] discuss and apply a Lagrangian dual for a minimum information problem
25
26 303 to show a way for obtaining more insight on the optimal solution.

4 AN ILLUSTRATIVE CASE-STUDY: ASSESSING SPATIAL DEPENDENCE OF POLITICAL VIOLENCE/ TERRORISM RISK IN INSURANCE UNDERWRITING

304 Given the flexibility and detail that the SRP method allows for when modelling dependence, we
305 regard it as of particular interest for application areas in which common simplifying assumptions,
306 such as bivariate normality, are not justified. Rather, different kinds of tail dependencies which
307 potentially induce extreme impact scenarios are prevalent. For these, we often assess and model
308 upper and lower tail dependence exclusively (similarly to testing the goodness of fit for asymmetric,
309 Archimedean copulas to historical data when available) given that e.g. joint large losses are typically
310 not observed together with joint large gains[40, 41].

311 As such, we consider (re-)insurance as an industry in which rigorous dependence modelling ap-
312 proaches are of particular interest. Due to the increasing complexity of (re-)insurance products,
313 new (holistic) modelling approaches, such as *dynamic financial analysis* (DFA) (a Monte Carlo
314 simulation-based method to model risks jointly), have become popular among actuaries to better
315 understand the risks an insurer underwrites [9]. For these new approaches, flexible and detailed
316 assessments of dependencies under a specific probability model are required. Exemplary for a DFA
317 application, Eling and Toplek [42] present how various parametric copulas can be used for stress-
318 testing an insurer's risk management strategies together with the implication on stakeholders, such

³In the context of expert judgement, an invariance approach to encoding information probabilistically is considered a main justification for maximum entropy methods North [38].

1
2
3
4 319 as regulators and rating agencies. The DFA model inputs, the perils (or risks) covered by an in-
5
6 320 surer, are informed by a *catastrophe model*. The components of catastrophe models are a hazard,
7
8 321 inventory, vulnerability and loss estimation module. The loss estimation output is usually an ex-
9
10 322 ceedance probability curve specifying probabilistically the severity levels of a certain hazard in a
11
12 323 region. Capturing relevant dependencies between severity levels is crucial for a more robust output.
13
14 324 See Grossi and Kunreuther [43] for an introduction to catastrophe models.

15
16 325 We have already established that a common challenge is lacking relevant historical data for quantify-
17
18 326 ing dependence relationships serving as model input. In actuarial risk assessment, *non-life* insurance
19
20 327 underwriting is particularly challenged. So called *low frequency-high severity* perils, natural and
21
22 328 man-made, are by definition not frequently observed but cannot be ignored. Therefore, we require
23
24 329 structured expert judgement to model their uncertainty. In this illustrative case-study, we apply
25
26 330 the SRP method to elicit and model the spatial dependence of the man-made peril of terrorism.
27
28 331 Terrorism attacks are not only often low frequency-high severity catastrophes but pose an additional
29
30 332 challenge due to intelligent adversaries which further inhibit the use of historical data. Better un-
31
32 333 derstanding the dependence between terrorism attacks' frequencies in different regions globally is
33
34 334 nevertheless key for an insurer to quantify and price this peril's risk when managing a portfolio of
35
36 335 (global) clients⁴.

336 4.1 Pricing terrorism risk in insurance

337 Traditionally, pricing of terrorism risk in insurance has not been evaluated from actuarial principles,
338
339 338 but rather covered by the balance of supply and demand in the insurance market together with some
340
341 339 less formal risk selection from site surveys [44]. Terrorism coverage (e.g. in the United States) had
342
343 340 been included in standard commercial insurance policies as an unnamed peril on all-risk commercial
344
345 341 and home owners coverages for property and contents [45]. More recent loss developments though
346
347 342 have highlighted the necessity of treating its risk assessment more rigorously. A major turning
348
349 343 point for dealing with terrorism risk in insurance was the attacks of September 11th, 2001 (9/11)
350
351 344 on the United States. The attacks incurred an estimated monetary loss up to 60 billion US dol-
352
353 345 lars, distributed among various lines of business, such as property insurance, business interruption
354
355 346 insurance and workers' compensation [46]. Globally, the worst 15 terrorist attacks in terms of casu-
356
357 347 alty numbers have occurred since 1982 with many more near-miss events [45]. Mathematically, the
358
359 348 relationship between the frequency of more recent attacks and their severity can be described by a
360
361 349 power law, i.e. attack severities that are orders of magnitude larger than the mean can be common

⁴As Woo [44] emphasises, we must not confuse quantifying terrorism risk with predicting a next attack. This is similar to natural catastrophes, such as earthquakes, for which we cannot determine the time, location and severity of the next event, but the aim is rather to evaluate the annual exceedance probability of loss, for instance to inform a property insurance portfolio.

[47]. The changing nature of its risk through an increasing number of frequencies and severities in multiple regions globally underlines the urgent need for improved assessment.

4.2 Expert judgement for adversarial risk

A specific aspect of assessing terrorism risk is the role of intelligent adversaries. Their impact is thus included in recent discussions on risk definitions [48, 49, 50, 51]. In fact, 9/11 led many researchers to propose modified risk definitions [52]. For instance, the triplet definition by Kaplan and Garrick [53] is extended to include adversaries in Garrick et al. [54] and Garrick [55] by considering the likelihood of a hazard as the conditional probability of a successful attack given that an attack is planned.

Models addressing adversarial risk are typically of game-theoretic nature [56, 57, 58] whereas the area of *adversarial risk analysis* comprises decision-analytic approaches combining traditional *probabilistic risk analysis* (PRA) methods with game theory [59, 60, 61, 62]. Nevertheless, there is some debate on (traditional) PRA's effectiveness for adversarial problems (see Ezell et al. [56] defending its usefulness and Brown and Cox Jr [63] and Cox Jr [64] arguing against it). A main argument against PRA approaches for adversaries is the dynamic attacker's decision rule for choosing a target as this choice might be based on the anticipated defender's assessment of targets' likelihoods. In other words, a defender's PRA might inform the attacker's choice and hence override its purpose as the previously most likely target has now zero probability of being attacked (closely related in terrorism risk analysis are decision on allocating defensive resources Bier [65]). Experts quantifying adversarial risk should therefore decompose their judgement in accordance with adversarial risk definitions, so that we understand experts' beliefs about attackers' choices. When doing so, assessments of an attack choice might be based on attackers' motivations, resources and capabilities together with defenders' vulnerabilities. In that way, expert judgement is used in the Probabilistic Terrorism Model by Risk Management Solutions Inc. (RMS⁵) for assessing likelihoods on target selection, capabilities of attack modes and an attack's overall likelihood. However, dependence between targets is neglected [66]. In other approaches, event trees are used to reason from an attacker's capabilities through a defender's countermeasures [67, 68]. In addition, several qualitative approaches for structuring the available knowledge on terrorists' objectives and motivations exist in the risk and decision analysis literature [69, 70, 54].

⁵RMS, founded at Stanford University in 1989, provides services in the area of catastrophe modelling for (re-)insurers.

4.3 Expert judgement for spatial dependence of terrorism attacks

Knowledge and beliefs on terrorists' motivations, resources and capabilities together with defender's vulnerabilities inform experts directly about the spatial dependence between attack frequencies.

Terrorist groups, such as the Irish Republican Army (IRA), Basque Separatist Group (ETA) or as well Hamas and the Palestine Liberation Organization (PLO), had and have specific geographical foci with a politically motivated attack purpose. Their goals are formulated and self-proclaimed as separatism or liberation. The attacks' geographical impact is identified straightforwardly. Based on the number of active terrorist groups per region plus their resources and capabilities relative to counter-measures, an expert assesses either positive or negative dependence. While positive dependence might not seem intuitive at first due to different local foci and typically a lack of collaboration between these groups, learning and encouragement by another groups' successes can still occur. Woo [71] regards learning of optimal behaviours beyond the own organisation as a main strength of some well-known terrorist groups. Other scenarios for positive dependence can be due to defenders' collaboration, joint counter-terrorism activities and sharing of intelligence resources.

In contrast to terrorists motivated by self-proclaimed liberatism and separatism, other groups derive their goals from religious ideology. These groups are often globally active. Their members are organised as multiple independent hubs with satellite cells. Al-Qaeda and the Islamic State of Iraq and Syria/the Levant (ISIS/ISIL) are typical examples of such network-based organisations [44, 72]. Models from swarm intelligence and statistical network analyses are used to evaluate the effectiveness of counter-terrorism measures and understand the attackers' capabilities. It is understood that organisations like Al-Qaeda and ISIS/ISIL are more resilient and capable of more severe attacks than (hierarchical) army-like structured groups [71]. For dependence assessments, understanding the global presence of members and sympathisers (potentially future recruits) together with the functioning of the network structure is crucial. For instance, scenarios of positive dependence can occur when a terrorist group obtains more power and resources to extend globally or when new attack types are used for which little intelligence or counter-measures exists. Scenarios of negative dependence might describe attackers' scarce resources, e.g. lacking financial support for regional hubs, so the target focus shifts towards a certain region. The latter also depends on vulnerabilities of target countries, desired attention through media or as well a planned revenge, e.g. for a country's military actions.

While these are only brief considerations for scenarios that can influence the assessment of dependence between the number of terrorist attacks in different regions (see Woo [71] for a more extensive discussion on regional and global terrorism), it shows the complexity of factors to be thought of. In this illustrative case-study, we focus on the geographical regions of Central Asia (CA) and Western




Figure 8: Regions of interest for dependence assessment.

413 Europe (WE) which are shown in Figure 8 (see the Appendix for a full list of the countries included
 414 per region).

415 4.4 Eliciting the marginal probabilities

416 Before eliciting dependence assessments from experts, we need to specify the marginal distributions
 417 for the variables of interest. Otherwise, the experts condition their judgements on different marginal
 418 probabilities and their assessments cannot be sensibly aggregated. The specification is done either
 419 through historical data (if available) or another, prior elicitation with a structured expert judgement
 420 method for univariate uncertainty, such as Quigley et al. [73], Gosling [74], Hanea et al. [75]. A
 421 structured elicitation for the marginal distributions is also encouraged when eliciting dependence
 422 only from one expert, i.e. without aggregation, as this mitigates potential biases of the marginals
 423 and ensures transparency[11].

424 In our case-study, the marginal distributions have been assessed by 16 experts⁶. The experts are
 425 involved in analysing and pricing the peril of terrorism and other armed conflict categories. They
 426 work for different (re-)insurers, catastrophe modellers and related service providers. The elicitation
 427 session was organised as part of the European Cooperation in Science and Technology, COST Action
 428 IS1304 - Expert Judgement Network, which aims at stimulating the emergence and spread of high
 429 quality evidence-based decision support approaches through structured expert judgement methods.
 430 The marginal distributions $F_X(x)$ and $F_Y(y)$ are defined as the number of terrorist attacks in Central
 431 Asia (x) and in Western Europe (y), both in 2017. We define a terrorist attack in accordance with
 432 common global data-bases on the topic (see START [76]). Thus, for an attack to be recorded as

⁶The 16 experts are from a first elicitation round from a currently ongoing study that aims to include more experts.

1
2
3
4 433 such there must be evidence of an intention to coerce, intimidate, or convey some other message to
5
6 434 a larger audience (or audiences) than the immediate victims. In this regard, any perpetrator group,
7
8 435 any weapon type (e.g. biological, chemical, explosive, firearms etc.), any attack type (e.g. armed
9
10 436 assault, bombing, facility/infrastructure attack, hostage taking etc.), any target apart from private
11
12 437 persons (i.e. business, infrastructure, military, educational/religious institutions etc.) is included.
13
14 438 We elicited $F_X(x)$ and $F_Y(y)$ through the so called *Classical Model* [73, 77]. Experts provide
15
16 439 various quantile assessments for a continuous quantity rather than point estimates. Usually (and
17
18 440 in our case), we elicit the 5th, 50th and 95th quantile. The experts answer two types of questions.
19
20 441 The first questions are about so called seed or calibration variables. For these, the true value is
21
22 442 known to the analyst but not the experts at the time of the elicitation (or they will be known later
23
24 443 and within the time frame of the study). The second question type is about the actual target value
25
26 444 or variable of interest, i.e. the uncertainties we intend to include in the model. Based on each
27
28 445 expert's assessments of the seed variables, the experts are aggregated. For that, two performance
29
30 446 measures are derived, the calibration and information score. Loosely, the calibration score measures
31
32 447 the statistical accuracy of the experts whose assessments are treated as statistical hypotheses. The
33
34 448 information score measures the assessments' concentrations relative to a background distribution.
35
36 449 Good expertise is shown by a high calibration and information score (see Cooke [77] for a more
37
38 450 detailed introduction). Figure 9 shows each experts' individual assessment for the target variables'
39
40 451 marginal distributions together with the aggregated assessments of equal weighting (EW) and the
41
42 452 classical method (DM global).

37 453 We observe in Figure 9 that the marginal distribution assessments are similar for both regions
38
39 454 whereas most of the experts provide narrow uncertainty bounds. The experts who are more uncertain
40
41 455 are so for both assessments. Hence, the performance-based and the equally weighted combination
42
43 456 show no major difference for either region. As commonly observed with the classical method, the
44
45 457 performance-based aggregation is more informative even if both combinations lead to similar median
46
47 458 assessments. **The full documentation, the elicitation protocol together with results and raw data for**
48
49 459 **the above elicitation can be found in Werner [78].**

460 **4.5 Applying the SRP method for quantifying spatial dependence of ter-** 461 **rorism risk**

462 Once the marginal distributions had been elicited, we proceeded with eliciting and modelling depen-
463
464 463 dence through the SRP method. This elicitation was done with a single expert who is a professional
465
466 464 in the area of terrorism catastrophe modelling within (re-)insurance (as well) and who subscribed to
467
468 465 the aggregated results for the marginal distributions. In total, we elicited six dependence judgements
469
470 466

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

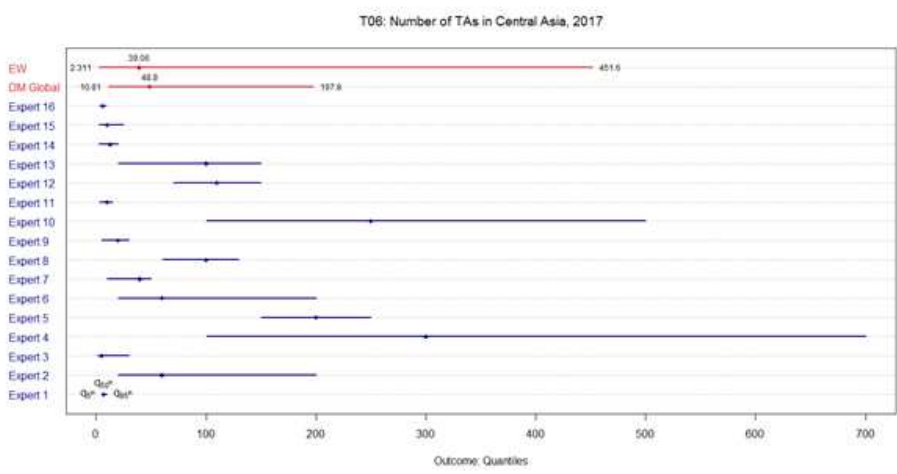
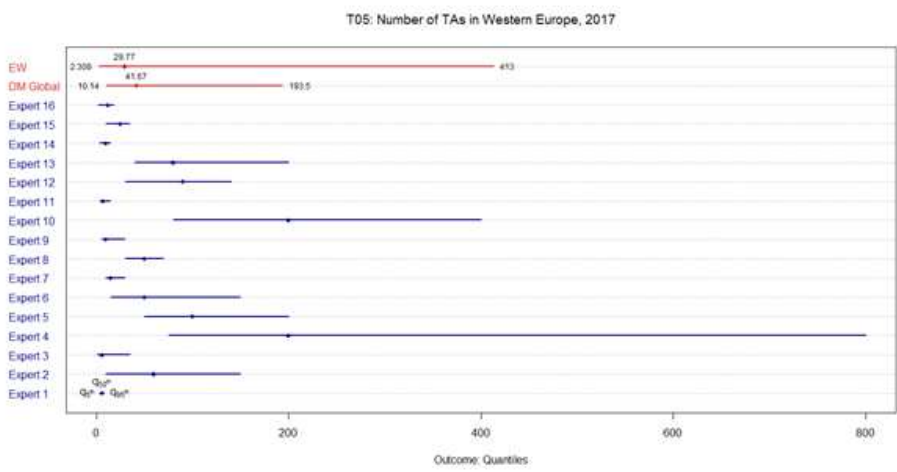


Figure 9: Outcome of eliciting the marginal distribution for each region.

Table I: Overview of dependence elicitation procedure and results.

	Framing	Conditional Probability	Assessment
(i)	"[...] more than 73 terrorist attacks in CA, what is your probability that we observe more than 62 terrorist attacks in WE?"	$P(Y > y_{0.5} X > x_{0.5})$	0.5
(ii)	"[...] more than 199 terrorist attacks in CA, what is your probability that we observe more than 62 terrorist attacks in WE?"	$P(Y > y_{0.5} X > x_{0.95})$	0.03
(iii)	"[...] more than 197 terrorist attacks in WE, what is your probability that we observe more than 73 terrorist attacks in CA?"	$P(X > x_{0.5} Y > y_{0.95})$	0.045
(iv)	"[...] more than 199 terrorist attacks in CA, what is your probability that we observe more than 197 terrorist attacks in WE?"	$P(Y > y_{0.95} X > x_{0.95})$	0.025
(v)	"[...] more than 199 terrorist attacks in CA, what is your probability that we observe more than 225 terrorist attacks in WE?"	$P(Y > y_{0.99} X > x_{0.95})$	0.04
(vi)	"[...] more than 225 terrorist attacks in WE, what is your probability that we observe more than 199 terrorist attacks in CA?"	$P(X > x_{0.95} Y > y_{0.99})$	0.01

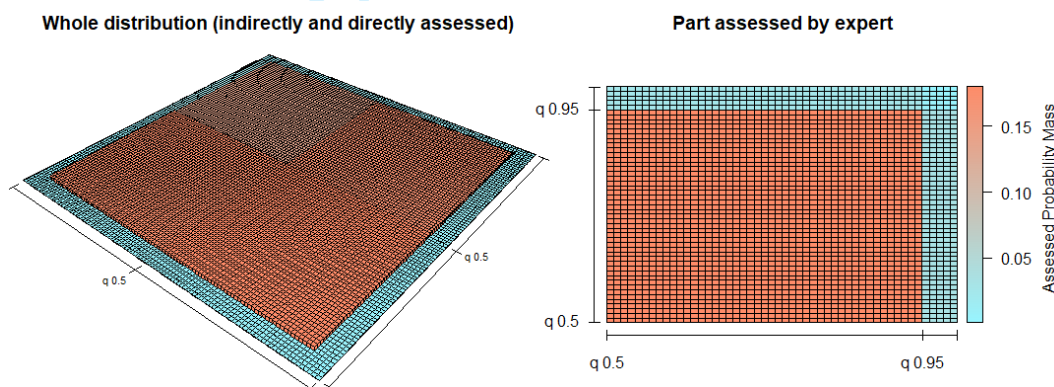


Figure 10: The experts joint distribution: overall (left) and assessed upper quadrant (right).

in addition to one further marginal assessment. The latter was required as we had not considered the 99th quantiles previously. As outlined in the initial exemplary procedure of section 2.2, we started by first eliciting an overall probability mass which was later partitioned to further explore the joint upper distribution tail. The first elicitation is therefore on the probability of the terrorist attack frequency in Western Europe (y) being above its 50th quantile, 62 attacks, given that we observe more than 73 attacks in Central Asia (x) (again the corresponding 50th quantile), both in the year 2017. All judgements were conditional probabilities given the expert's familiarity with its interpretation. Table I summarises the dependence assessments by showing the results together with the framing of the questions.

As second part of the SRP method, we then modelled the overall joint distribution for the spatial dependence through solving the minimum information minimisation problem (section 3) based on the above assessments. The result can be seen in Figure 10.

We see that the expert's distribution indicates a slight negative dependence relationship be-

1
2
3
4 479 tween the spatial terrorism risk of both regions which is however close to independence. This is
5 480 particularly driven by the first assessment being equal to 0.5 which indicates independence for a
6 481 broad area of the joint distribution. In more detail, the difference between assessment *ii.*) and *iii.*)
7 482 shows that in the joint tail, the expert assesses that an extreme year in terms of number of attacks
8 483 for WE affects CA more than vice versa. The slight negative dependence (close to independence)
9 484 corresponds to the expert's rationale which has been formally facilitated in order to support the
10 485 expert with structuring his/her knowledge about the spatial dependence between both regions. For
11 486 that, we used a conditional scenario mapping method [34]. In addition to mitigating some prevalent
12 487 cognitive fallacies of assessing dependence, such as the *confusion of the inverse* or *confusing joint*
13 488 *and conditional probabilities* (see also Werner et al. [11] for an overview), this method allows for
14 489 considering and reflecting explicitly which scenarios affect the probability spaces of both regions (in
15 490 a conditional sense). Scenarios are defined as "sequences that link triggering events to specified con-
16 491 sequences (or final states) through intermediate conditions" [34]. For the example shown in Figure
17 492 11, the expert first reasoned through backwards logic, i.e. starting from the specified consequence,
18 493 about observing more than 199 in Central Asia until the end of 2017 (α_5). Then, based on the
19 494 initiating events that might cause Central Asia to experience more than 199 attacks and which are
20 495 (at least partly) observable today, the expert reasoned (in forward logic) how these same initiating
21 496 events affect the development of the number of terrorist attacks in Western Europe until the end of
22 497 2017. Based on the the number and plausibility of these conditional scenarios causing more than 255
23 498 attacks (again α_5), the expert could then make a dependence assessment in a more informed and
24 499 confident manner. Werner et al. [34] presents the structured process of generating such conditional
25 500 scenarios in more detail.

26 501 As can be seen in Figure 11, the expert considers both regions to be slightly negatively dependent
27 502 (close to independence) due to the consideration that the active terrorist groups in both regions are
28 503 different. In Central Asia, local separatists have political and regional motivations while in Western
29 504 Europe mainly islamist groups are prevalent despite e.g. Russia's military involvement in the Middle
30 505 East. Furthermore, the expert considers both regions to be different with regards to their vulnera-
31 506 bility given not only the types of active terrorist groups but also the varying counter-terrorism and
32 507 intelligence capabilities which drive the negative dependence relationship.

33 508 Before concluding this illustrative example, a first remark is that for quantifying the spatial depen-
34 509 dence of terrorism attacks the definition used in this example is rather broad by including all attack
35 510 types. Thus, the consideration of specific attack types might have very particular effects on the
36 511 geographical interdependencies. As such, of growing interest in the adversarial risk literature have
37 512 been biological attacks [56] and cyber attacks [79]. For these, it can be informative to assess the




Figure 11: Unconditional and conditional scenarios for assessments.

1
2
3
4 513 dependence between variables of interest, such as casualties or monetary losses.
5
6 514 Further, we understand that an elicitation considering more explicitly the geographical interdepen-
7
8 515 dencies of critical infrastructure can be informative for insurers, for instance when offering business
9
10 516 interruption coverage. Our method could hence build upon some modelling approaches that have
11
12 517 ranked the susceptibility of critical infrastructures targeted by attacks [80].
13
14 518 Lastly, we acknowledge the inherent difficulties particular when considering attacks, such as 9/11,
15
16 519 which some might title "black swans". For dependencies, the term "perfect storms" appeared (see
17
18 520 Paté-Cornell [81] for a discussion on the use of these terms in risk analysis and management).
19
20 521 However, even for such events, structured assessment through experts can be informative and it is
21
22 522 interesting that e.g. Zelikow (as director of the 9/11 Commission) called the misreading of precur-
23
24 523 sors to these events as "failure of imagination" given that air-planes had been used before as weapon
25
26 524 and the World Trade Center in New York had been targeted already in 1993 [81].

27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59 60

5 DISCUSSION AND CONCLUSIONS

525 When using expert judgement for assessing dependence, there is a trade-off between easing the
526
527 assessment burden for experts and sufficiently capturing a real-world phenomenon of interest in our
528
529 model [10]. Therefore, we have presented an elicitation method that aims to satisfy a decision-
530
531 maker's desired level of detail for a model whereas the procedure for eliciting dependence from
532
533 experts provides an intuitive way of assessing even detailed dependence information (such as extreme
534
535 parts of a joint distribution) while avoiding infeasible and inconsistent assessments. We argued that
536
537 for the decision-maker a non-parametric setting of modelling multivariate uncertainties is more
538
539 desirable and therefore we addressed the potential assessment issues of under- and overspecification.
540
541 Concluding on the application shown in this paper, we note that in future research more applications
542
543 are desirable to explore how the SRP method performs and obtain insights on potential modifications
544
545 like alternative ways of framing the judgements, the implication of restricted feasibility ranges, or the
546
547 elicitation of different forms (other than conditional probabilities). For example, as an alternative
548
549 to eliciting quantile-based assessments, we can elicit conditional expectations. This follows from the
550
551 discussion of Werner et al. [10] on modelling and elicitation strategies that are determined by the
552
553 choice of considering influencing factors of dependence relationships explicitly or implicitly. The
554
555 latter is similar to PI methods which aim at satisfying reasonable conditions of a model output due
556
557 to its easier understanding and quantification. This is of particular interest when we cannot observe
558
559 (and hence elicit) our variables of interest directly. Bedford et al. [19] show an elicitation procedure
560
561 and minimum information modelling method for expectations on the whole joint distribution. Hence,
562
563 considering its elaboration based on our method could allow for a more detailed specification of

1
2
3
4 545 multivariate uncertainty for non-observable model input parameters. In the actuarial context of
5
6 546 section 4, we might ask experts to assess the conditional expectation for a risk measure, such as
7
8 547 *probable maximum loss* (see Grossi and Kunreuther [43]), which can be used (partly) as model
9
10 548 output, whereas we assess dependence through PI on the function generating it. Similarly, our
11
12 549 method can be used, either through quantile-based assessments or modifications, in other sectors
13
14 550 in which understanding and quantifying tail risk is becoming of growing interest, such as financial
15
16 551 decision-making on asset allocation [82].

17 ACKNOWLEDGEMENTS

20 APPENDIX

23 Proof for Proposition 1:

24 552 Suppose we are given values $u_0 = 0 < u_1 < \dots < u_n < 1 = u_{n+1}$, and $v_0 = 0 < v_1 < \dots < v_m <$
25
26 553 $1 = v_{m+1}$ (where $n, m > 0$), $0 < p, q < 1$, with p different to the u_i and q different to the v_j . Then
27
28 554 a copula distribution on $QP(u, v)$ can be refined to a copula distribution on $QP(u, v; p, q)$.
29
30 555

31 556 *Proof.* In order to prove proposition 1, we divide the set $QP(u, v)$ into four subsets:
32
33 557

- 34
35 558 1. $A(p, q)$ has a single element which is the rectangle of $QP(u, v)$ containing the point (p, q) .
- 36
37 559 2. $U(p, q)$ is the set of rectangles in $QP(u, v)$ that overlap the line $v = q$, except the one in
38
39 560 $A(p, q)$.
- 40
41 561 3. $V(p, q)$ is the set of rectangles in $QP(u, v)$ that overlap the line $u = p$, except the one in
42
43 562 $A(p, q)$.
- 44
45 563 4. $B(p, q)$ is the set of all rectangles in $QP(u, v)$ that are not in $A(p, q)$, $B(p, q)$, or $V(p, q)$.

46
47 564 Define also $A^*(p, q)$ to be the rectangles in $QP(u, v; p, q)$ which are sub-rectangles of $A(p, q)$, and
48
49 565 define U^* , V^* and B^* similarly.

50
51 566 Note that $B^*(p, q) = B(p, q)$, that is, the rectangles in $B(u, v)$ do not get subdivided by the lines
52
53 567 $u = p$, $v = q$. Rectangles in U^* are obtained by dividing rectangles in U by the line $v = q$, and
54
55 568 rectangles in V^* are obtained by dividing rectangles in V by the line $u = p$.

56
57 569 We now define the refined copula distribution on $QP(u, v; p, q)$.

58
59 570 Let $\alpha = (p - u_i)/(u_{i+1} - u_i)$, and $\beta = (q - v_j)/(v_{j+1} - v_j)$. We specify how to define the refined
60
571 copula distribution as follows:




Figure 12: Different set of rectangles in $QP(u, v)$.

1. For the rectangles in A^* , the lower left sub-rectangle is allocated $\alpha\beta$ of the mass of A , the lower right one gets proportion $(1 - \alpha)\beta$, the upper left one gets proportion $\alpha(1 - \beta)$, and the upper right one gets proportion $(1 - \alpha)(1 - \beta)$.
2. Each rectangle in U is subdivided into two sub-rectangles in U^* by the line $v = q$, and the lower sub-rectangle is allocated proportion β of its mass and the upper sub-rectangle is allocated proportion $(1 - \beta)$ of the mass.
3. Each rectangle in V is subdivided into two sub-rectangles in V^* by the line $u = p$, whereas the left sub-rectangle is allocated proportion α of its mass and the upper sub-rectangle is allocated proportion $(1 - \alpha)$ of its mass.
4. Any rectangle in $B^*(p, q) = B(p, q)$ is assigned the same probability as it was in in the copula distribution on $QP(u, v)$.

This allocation of probabilities to the rectangles of $QP(u, v; p, q)$ adds to 1, while it is straightforward to check that it is a copula distribution. \square

Regions of interest in illustrative case-study (section 4):

- Central Asia: Armenia, Azerbaijan, Belarus, Georgia, Kazakhstan, Kyrgyzstan, Russia, Tajikistan, Turkmenistan, Ukraine, Uzbekistan.
- Western Europe: Austria, Belgium, Denmark, Finland, France, Germany, Iceland, Ireland, Italy, Luxembourg, Malta, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom.

REFERENCES

- 591 [1] James E Smith and Detlof Von Winterfeldt. Anniversary article: decision analysis in manage-
592 ment science. *Management Science*, 50(5):561–574, 2004.
- 593 [2] Robert L Winkler. State of the art: research directions in decision making under uncertainty.
594 *Decision Sciences*, 13(4):517–533, 1982.
- 595 [3] Fabrizio Durante and Carlo Sempi. *Principles of copula theory*. CRC Press, Boca Raton, 2015.
- 596 [4] Anca Hanea, Oswaldo Morales Napoles, and Dan Ababei. Non-parametric bayesian networks:
597 Improving theory and reviewing applications. *Reliability Engineering & System Safety*, 144:
598 265–284, 2015.
- 599 [5] Alexander J McNeil, Rüdiger Frey, and Paul Embrechts. *Quantitative risk management: Con-
600 cepts, techniques and tools*. Princeton University Press, USA, 2015.
- 601 [6] Harry Joe. *Dependence modeling with copulas*. CRC Press, Boca Raton, 2014.
- 602 [7] Christian Genest, Michel Gendron, and Michaël Bourdeau-Brien. The advent of copulas in
603 finance. *The European Journal of Finance*, 15(7-8):609–618, 2009.
- 604 [8] Dorota Kurowicka and Roger M Cooke. *Uncertainty analysis with high dimensional dependence
605 modelling*. John Wiley & Sons, Chichester, 2006.
- 606 [9] P Embrechts, A McNeil, and D Straumann. Correlation and dependence in risk management:
607 properties and pitfalls. *Risk management: value at risk and beyond*, pages 176–223, 2002.
- 608 [10] Christoph Werner, Tim Bedford, Roger M Cooke, Anca M Hanea, and Oswaldo Morales-
609 Nápoles. Expert judgement for dependence in probabilistic modelling: a systematic literature
610 review and future research directions. *European Journal of Operational Research*, 258(3):801–
611 819, 2017.
- 612 [11] Christoph Werner, Anca M Hanea, and Oswaldo Morales-Nápoles. Eliciting multivariate un-
613 certainty from experts: Considerations and approaches along the expert judgement process.
614 In Luis C Dias, Alec Morton, and John Quigley, editors, *Elicitation: The science and art of
615 structuring judgement*, volume 261, chapter 8, pages 171–210. Springer International Series in
616 Operations Research and Management Science, New York, 2018.
- 617 [12] Roger M Cooke and Louis H J Goossens. Procedures guide for structured expert judgment
618 project report eur 18820en. , European Commission Euratom, 2000.

- 1
2
3
4 619 [13] Roger M Cooke and Louis H J Goossens. Procedures guide for structural expert judgement in
5 accident consequence modelling. *Radiation Protection Dosimetry*, 90(3):303–311, 2000.
6
7
8 621 [14] Bruno De Finetti. Logical foundations and measurement of subjective probability. *Acta Psy-*
9 *chologica*, 34:129–145, 1970.
10
11
12 623 [15] Leonard Jimmie Savage. *The foundations of statistics*. John Wiley and Sons, New York, 1954.
13
14 624 [16] Frank P Ramsey. Truth and probability. *The foundations of mathematics and other logical*
15 *essays*, pages 156–198, 1931.
16
17
18 626 [17] Edwin T Jaynes. *Probability theory: The logic of science*. Cambridge University Press, Cam-
19 *bridge*, 2003.
20
21
22 628 [18] Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*.
23 Morgan Kaufmann, 1988.
24
25
26 630 [19] Tim Bedford, Kevin J Wilson, and Alireza Daneshkhah. Assessing parameter uncertainty on
27 coupled models using minimum information methods. *Reliability Engineering & System Safety*,
28 125:3–12, 2014.
29
30
31
32 633 [20] Bernd Kraan and Tim Bedford. Probabilistic inversion of expert judgments in the quantification
33 of model uncertainty. *Management Science*, 51(6):995–1006, 2005.
34
35
36 635 [21] Adrianus Marinus Hendrikus Meeuwissen and Tim Bedford. Minimally informative distribu-
37 tions with given rank correlation for use in uncertainty analysis. *Journal of Statistical Compu-*
38 *tation and Simulation*, 57(1-4):143–174, 1997.
39
40
41 638 [22] Roger M Cooke. Parameter fitting for uncertain models: modelling uncertainty in small models.
42 *Reliability Engineering & System Safety*, 44(1):89–102, 1994.
43
44
45 640 [23] Tim Bedford, Alireza Daneshkhah, and Kevin J Wilson. Approximate uncertainty modeling in
46 risk analysis with vine copulas. *Risk Analysis*, 36:792–815, 2016.
47
48
49 642 [24] Tim Bedford and Roger M Cooke. Vines: A new graphical model for dependent random
50 variables. *Annals of Statistics*, pages 1031–1068, 2002.
51
52
53 644 [25] Luis V Montiel and J Eric Bickel. Approximating joint probability distributions given partial
54 information. *Decision Analysis*, 10(1):26–41, 2013.
55
56
57 646 [26] Luis V Montiel and J Eric Bickel. A simulation-based approach to decision making with partial
58 information. *Decision Analysis*, 9(4):329–347, 2012.
59
60

- 1
2
3
4 648 [27] Ali E Abbas. Entropy methods for joint distributions in decision analysis. *IEEE Transactions*
5 649 *on Engineering Management*, 53(1):146–159, 2006.
- 6
7
8 650 [28] J Eric Bickel and James E Smith. Optimal sequential exploration: A binary learning model.
9 651 *Decision Analysis*, 3(1):16–32, 2006.
- 10
11
12 652 [29] Thomas Nagler and Claudia Czado. Evading the curse of dimensionality in nonparametric
13 653 density estimation with simplified vine copulas. *Journal of Multivariate Analysis*, 151:69–89,
14 654 2016.
- 15
16
17
18 655 [30] Robert J Vanderbei. *Linear programming*. Springer, 2015.
- 19
20 656 [31] Roger B Nelsen. *An introduction to copulas*. Springer Science & Business Media, 2007.
- 21
22 657 [32] C S Spetzler and C A S Stael von Holstein. Exceptional paper: Probability encoding in decision
23 658 analysis. *Management Science*, 22(3):340–358, 1975.
- 24
25
26 659 [33] Bansal S and A Palley. Is it better to elicit quantile or probability judgments? a comparison
27 660 of direct and calibrated procedures for estimating a continuous distribution. *Kelley School of*
28 661 *Business Research Paper*, pages 17–44, 2017.
- 29
30
31
32 662 [34] Christoph Werner, Tim Bedford, and John Quigley. Mapping conditional scenarios for knowl-
33 663 edge structuring in (tail) dependence elicitation. *Journal of the Operational Research Society*,
34 664 under review.
- 35
36
37 665 [35] Benjamin Margolin Rottman and Reid Hastie. Reasoning about causal relationships: Inferences
38 666 on causal networks. *Psychological bulletin*, 140(1):1–31, 2013.
- 39
40
41 667 [36] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of*
42 668 *mathematical statistics*, 22(1):79–86, 1951.
- 43
44
45 669 [37] Tim Bedford and Kevin J Wilson. On the construction of minimum information bivariate
46 670 copula families. *Annals of the Institute of Statistical Mathematics*, 66(4):703–723, 2014.
- 47
48
49 671 [38] D Warner North. *The Invariance Approach to the Probabilistic Encoding of Information*. PhD
50 672 thesis, Stanford University, Stanford, USA, 1970.
- 51
52
53 673 [39] Claude E Shannon. A mathematical theory of communication. *Bell System Technical Journal*,
54 674 27(3):379–423, 1948.
- 55
56
57 675 [40] Carolyn Kousky and Roger M Cooke. The unholy trinity: fat tails, tail dependence, and
58 676 micro-correlations. *Resources for the Future Discussion Paper*, 9(36):1–36, 2009.

- 1
2
3
4 677 [41] Oswaldo Morales-Nápoles, Dominik Paprotny, Daniël Worm, Linda Abspoel-Bukman, and Wim
5 678 Courage. Characterization of precipitation through copulas and expert judgement for risk
6 679 assessment of infrastructure. *ASCE-ASME Journal of Risk and Uncertainty in Engineering*
7 680 *Systems, Part A: Civil Engineering*, 3(4):3–18, 2017.
- 10 681 [42] Martin Eling and Denis Toplek. Modeling and management of nonlinear dependencies–copulas
11 682 in dynamic financial analysis. *Journal of Risk and Insurance*, 76(3):651–681, 2009.
- 14 683 [43] Patricia Grossi and Howard Kunreuther. *Catastrophe modeling: a new approach to managing*
15 684 *risk*. Springer Science & Business Media, New York, 2005.
- 18 685 [44] Gordon Woo. Quantitative terrorism risk assessment. *The Journal of Risk Finance*, 4(1):7–14,
19 686 2002.
- 22 687 [45] Erwann Michel-Kerjan and Burkhard Pedell. How does the corporate world cope with mega-
23 688 terrorism? puzzling evidence from terrorism insurance markets. *Journal of Applied Corporate*
24 689 *Finance*, 18(4):61–75, 2006.
- 27 690 [46] Jozef De Mey. The aftermath of september 11: The impact on and systemic risk to the insurance
28 691 industry. *The Geneva Papers on Risk and Insurance Issues and Practice*, 28(1):65–70, 2003.
- 31 692 [47] Aaron Clauset, Maxwell Young, and Kristian Skrede Gleditsch. On the frequency of severe
32 693 terrorist events. *Journal of Conflict Resolution*, 51(1):58–87, 2007.
- 35 694 [48] Terje Aven and Seth Guikema. On the concept and definition of terrorism risk. *Risk Analysis*,
36 695 35(12):2162–2171, 2015.
- 39 696 [49] Terje Aven and Bodil S Krohn. A new perspective on how to understand, assess and manage
40 697 risk and the unforeseen. *Reliability Engineering & System Safety*, 121:1–10, 2014.
- 43 698 [50] Terje Aven. The risk concept-historical and recent development trends. *Reliability Engineering*
44 699 *& System Safety*, 99:33–44, 2012.
- 47 700 [51] Louis Anthony Cox Jr. Some limitations of 'risk= threat× vulnerability× consequence' for risk
48 701 analysis of terrorist attacks. *Risk Analysis*, 28(6):1749–1761, 2008.
- 51 702 [52] Yacov Y Haimes. On the complex definition of risk: A systems-based approach. *Risk Analysis*,
52 703 29(12):1647–1654, 2009.
- 55 704 [53] Stanley Kaplan and B John Garrick. On the quantitative definition of risk. *Risk Analysis*, 1
56 705 (1):11–27, 1981.

- 1
2
3
4 706 [54] B John Garrick, James E Hall, Max Kilger, John C McDonald, Tara O’Toole, Peter S Probst,
5 Elizabeth Rindskopf Parker, Robert Rosenthal, Alvin W Trivelpiece, Lee A Van Arsdale, et al.
6 707
7 708 Confronting the risks of terrorism: making the right decisions. *Reliability Engineering & System*
8 *Safety*, 86(2):129–176, 2004.
9 709
- 10
11 710 [55] John B Garrick. Perspectives on the use of risk assessment to address terrorism. *Risk Analysis*,
12 22(3):421–423, 2002.
13 711
- 14
15 712 [56] Barry Charles Ezell, Steven P Bennett, Detlof Von Winterfeldt, John Sokolowski, and Andrew J
16 Collins. Probabilistic risk analysis and terrorism risk. *Risk Analysis*, 30(4):575–589, 2010.
17 713
- 18
19 714 [57] Louis Anthony Tony Cox Jr. Game theory and risk analysis. *Risk Analysis*, 29(8):1062–1068,
20 2009.
21 715
- 22
23 716 [58] Elisabeth Paté-Cornell and Seth Guikema. Probabilistic modeling of terrorist threats: A sys-
24 tems analysis approach to setting priorities among countermeasures. *Military Operations Re-*
25 *search*, 7(4):5–23, 2002.
26 717
27 718
- 28
29 719 [59] Eduardo G Quijano, David Ríos Insua, and Javier Cano. Critical networked infrastructure
30 protection from adversaries. *Reliability Engineering & System Safety*, 2016. doi: 10.1016/j.
31 ress.2016.10.015.
32 720
33 721
- 34
35 722 [60] Jesus Rios and David Rios Insua. Adversarial risk analysis for counterterrorism modeling. *Risk*
36 *Analysis*, 32(5):894–915, 2012.
37 723
- 38
39 724 [61] Seth D Guikema and Terje Aven. Assessing risk from intelligent attacks: A perspective on
40 approaches. *Reliability Engineering and System Safety*, 15(8):478–483, 2010.
41 725
- 42
43 726 [62] David Rios Insua, Jesus Rios, and David Banks. Adversarial risk analysis. *Journal of the*
44 *American Statistical Association*, 104(486):841–854, 2009.
45 727
- 46
47 728 [63] Gerald G Brown and Louis Anthony Tony Cox Jr. How probabilistic risk assessment can
48 mislead terrorism risk analysts. *Risk Analysis*, 31(2):196–204, 2011.
49 729
- 50
51 730 [64] Louis Anthony Cox Jr. Improving risk-based decision making for terrorism applications. *Risk*
52 *Analysis*, 29(3):336–341, 2009.
53 731
- 54
55 732 [65] Vicki M Bier. Choosing what to protect. *Risk Analysis*, 27(3):607–620, 2007.
- 56
57 733 [66] Henry H Willis. Guiding resource allocations based on terrorism risk. *Risk analysis*, 27(3):
58 597–606, 2007.
59 734
60

- 1
2
3
4 735 [67] Heather Rosoff and Detlof Von Winterfeldt. A risk and economic analysis of dirty bomb attacks
5 736 on the ports of los angeles and long beach. *Risk Analysis*, 27(3):533–546, 2007.
- 6
7
8 737 [68] Detlof Von Winterfeldt and Terrence M O’Sullivan. Should we protect commercial airplanes
9 738 against surface-to-air missile attacks by terrorists? *Decision Analysis*, 3(2):63–75, 2006.
- 10
11
12 739 [69] Johannes Siebert, Detlof von Winterfeldt, and Richard S John. Identifying and structuring the
13 740 objectives of the islamic state of iraq and the levant (isil) and its followers. *Decision Analysis*,
14 741 13(1):26–50, 2015.
- 15
16
17 742 [70] Gregory L Keeney and Detlof Von Winterfeldt. Identifying and structuring the objectives of
18 743 terrorists. *Risk Analysis*, 30(12):1803–1816, 2010.
- 19
20
21 744 [71] Gordon Woo. *Calculating catastrophe*. World Scientific, 2011.
- 22
23
24 745 [72] Gordon Woo. Understanding the principles of terrorism risk moeling from the attack in west-
25 746 minster. , Risk Management Solutions, London, UK, Discussion paper, 2017.
- 26
27
28 747 [73] John Quigley, Abigail Colson, Willy Aspinall, and Roger Cooke. Elicitation in the classical
29 748 model. In Luis C Dias, Alec Morton, and John Quigley, editors, *Elicitation: The science and*
30 749 *art of structuring judgement*, volume 261, chapter 2, pages 15–36. Springer International Series
31 750 in Operations Research and Management Science, New York, 2018.
- 32
33
34 751 [74] John Paul Gosling. Shelf: the sheffield elicitation framework. In Luis C Dias, Alec Morton, and
35 752 John Quigley, editors, *Elicitation: The science and art of structuring judgement*, volume 261,
36 753 chapter 4, pages 61–93. Springer International Series in Operations Research and Management
37 754 Science, New York, 2018.
- 38
39
40
41 755 [75] Anca Maria Hanea, Mark Burgman, and Victoria Hemming. Idea for uncertainty quantification.
42 756 In Luis C Dias, Alec Morton, and John Quigley, editors, *Elicitation: The science and art of*
43 757 *structuring judgement*, volume 261, chapter 5, pages 95–117. Springer International Series in
44 758 Operations Research and Management Science, New York, 2018.
- 45
46
47
48
49 759 [76] START. Global terrorism database, 2016. data retrieved from: [https://www.start.umd.edu/](https://www.start.umd.edu/gtd)
50 760 [gtd](https://www.start.umd.edu/gtd).
- 51
52
53 761 [77] Roger M Cooke. *Experts in uncertainty: opinion and subjective probability in science*. Oxford
54 762 University Press, 1991.
- 55
56
57 763 [78] Christoph Werner. Geopolitical risk assessment: A structured expert judgement elic-
58 764 itation - elicitation protocol and results, 2018. URL [http://dx.doi.org/10.15129/](http://dx.doi.org/10.15129/a0029a02-4283-4bc8-a671-ee5dbe680dfa)
59 765 [a0029a02-4283-4bc8-a671-ee5dbe680dfa](http://dx.doi.org/10.15129/a0029a02-4283-4bc8-a671-ee5dbe680dfa).

- 1
2
3
4 766 [79] Jinsoo Shin, Hanseong Son, Gyunyoung Heo, et al. Development of a cyber security risk model
5 767 using bayesian networks. *Reliability Engineering & System Safety*, 134:208–217, 2015.
6
7
8 768 [80] Sean Albert Patterson and George E Apostolakis. Identification of critical locations across
9 769 multiple infrastructures for terrorist actions. *Reliability Engineering & System Safety*, 92(9):
10 770 1183–1203, 2007.
11
12
13 771 [81] Elisabeth Paté-Cornell. On 'black swans' and 'perfect storms': risk analysis and management
14 772 when statistics are not enough. *Risk Analysis*, 32(11):1823–1833, 2012.
15
16
17 773 [82] George P Gao, Xiaomeng Lu, and Zhaogang Song. Tail risk concerns everywhere. *Management*
18 774 *Science*, 2018.
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

For Peer Review