Department of Mathematics and Systems Analysis

Computational models for adversarial risk analysis and probabilistic scenario planning

Juho Roponen

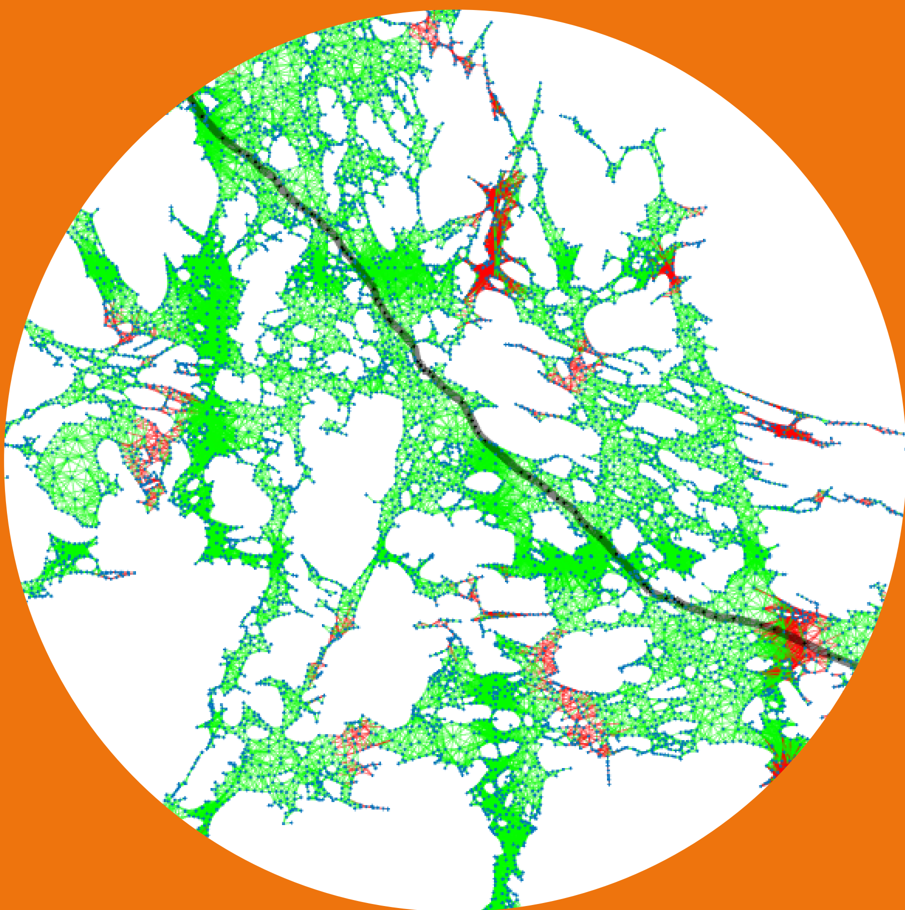

Aalto University

DOCTORAL THESES

Aalto University publication series
DOCTORAL THESES 67/2023

Computational models for adversarial risk analysis and probabilistic scenario planning

Juho Roponen

A doctoral thesis completed for the degree of Doctor of Science (Technology) to be defended, with the permission of the Aalto University School of Science, at a public examination held at the lecture hall H304 (Otakaari 1, Espoo, Finland) of the school on 9 June 2023 at 12 noon.

Aalto University
School of Science
Department of Mathematics and Systems Analysis

Supervising professor
Professor Ahti Salo, Aalto University, Finland

Thesis advisor
Professor Ahti Salo, Aalto University, Finland

Preliminary examiners
Professor Ali Abbas, University of Southern California, United States of America
Professor Aron Larsson, Stockholm University and Mid-Sweden University, Sweden

Opponent
Professor Ali Abbas, University of Southern California, United States of America

Aalto University publication series
DOCTORAL THESES 67/2023

ISBN 978-952-64-1260-3 (printed)
ISBN 978-952-64-1261-0 (pdf)
ISSN 1799-4934 (printed)
ISSN 1799-4942 (pdf)
http://urn.fi/URN:ISBN:978-952-64-1261-0

Unigrafia Oy
Helsinki 2023

Finland

NORDIC SWAN ECOLABEL
Printed matter
4041-0619

Aalto University

Aalto University, P.O. Box 11000, FI-00076 Aalto www.aalto.fi

Abstract

Author
Juho Roponen

Name of the doctoral thesis
Computational models for adversarial risk analysis and probabilistic scenario planning

Publisher School of Science

Unit Department of Mathematics and Systems Analysis

Series Aalto University publication series DOCTORAL THESES 67/2023

Field of research Systems and Operations Research

Manuscript submitted 26 April 2023 **Date of the defence** 9 June 2023

Permission for public defence granted (date) 12 April 2023 **Language** English

☐ **Monograph** ☒ **Article thesis** ☐ **Essay thesis**

Abstract
People need to make decisions under uncertainty. Both in corporate and public governance, in addition to uncertainty, the decisions can have high costs and far-reaching consequences. Thus, choosing a good decision alternative, or at least avoiding the inferior ones, is crucial. Two sources of uncertainty are especially prevalent in these decision problems: human activity and long planning horizons.

In this dissertation, methods for addressing uncertainties arising from both these sources are developed. By quantifying these uncertainties as probability distributions and preferences over outcomes as utility functions, a well-defined mathematical decision problem can be constructed and then solved using optimization techniques.

First, methods for adversarial risk analysis are developed to model the decision processes of adversarial actors who deliberately try to advance their own interests. The proposed methods facilitate systematic probabilistic analyses with limited knowledge about the adversary's preferences and their available information. This can be especially useful when the exact way the adversary analyzes the situation is difficult to assess or when their goals are deliberately hidden, as is often the case when analyzing military combat or security problems. The dissertation also demonstrates how combat modeling and simulation tools can be applied in adversarial risk analysis. This expands the types of analyses these tools can be used for, making it possible to answer questions such as, how the adversary's actions are impacted by changing circumstances, or how the outcomes of individual battles impact the larger strategic situation.

Second, a new probabilistic cross-impact analysis model is developed to quantify uncertainties associated with future scenarios based on information elicited from subject matter experts. Two different computational approaches are presented for analyzing the elicited cross-impact statements. One takes information about upper and lower bounds on probabilities and then calculates upper and lower bounds on system risk or utility. The other takes the best estimates about probabilities of specific uncertainty factors and their interactions and constructs a joint probability distribution and a Bayesian network. These approaches can be useful when probability information based on statistics or simulations is not available, for example when results need to be produced quickly or the uncertainties are associated with relatively far-off future events or human activity.

Keywords risk analysis, game theory, decision analysis, scenario analysis, cross-impact analysis

ISBN (printed) 978-952-64-1260-3 **ISBN (pdf)** 978-952-64-1261-0

ISSN (printed) 1799-4934 **ISSN (pdf)** 1799-4942

Location of publisher Helsinki **Location of printing** Helsinki **Year** 2023

Pages 122 **urn** http://urn.fi/URN:ISBN:978-952-64-1261-0

Aalto-yliopisto

Aalto-yliopisto, PL 11000, 00076 Aalto www.aalto.fi

Tiivistelmä

Tekijä
Juho Roponen

Väitöskirjan nimi
Laskennallisia malleja vastakkainasettelulliseen riskianalyysiin ja todennäköisyyspohjaiseen skenaariosuunnitteluun

Julkaisija Perustieteiden korkeakoulu

Yksikkö Matematiikan ja systeemianalyysin laitos

Sarja Aalto University publication series DOCTORAL THESES 67/2023

Tutkimusala Systeemi- ja operaatiotutkimus

Käsikirjoituksen pvm 26.04.2023 **Väitöspäivä** 09.06.2023

Väittelyluvan myöntämispäivä 12.04.2023 **Kieli** Englanti

☐ **Monografia** ☒ **Artikkeliväitöskirja** ☐ **Esseeväitöskirja**

Tiivistelmä

Ihmiset joutuvat väistämättä tekemään päätöksiä epävarmuuden vallitessa. Yritysjohdossa ja julkishallinnossa moniin päätöksiin liittyy epävarmuuden lisäksi myös paljon kustannuksia ja kauaskantoisia seurauksia. Siksi hyvän päätösvaihtoehdon löytäminen, tai ainakin huonojen välttäminen, on ensiarvoisen tärkeää. Näissä konteksteissa epävarmuudet liittyvät yleisimmin joko ihmisten toiminnan tai tulevaisuuden heikkoon ennakoitavuuteen.

Tässä väitöskirjassa kehitetään menetelmiä kummankinlaisten epävarmuuksien käsittelyyn. Kuvaamalla epävarmuuksia todennäköisyysjakaumilla ja seurauksien haluttavuutta hyötyfunktioilla on mahdollista rakentaa hyvin määritelty matemaattinen päätösongelma ja ratkaista se optimointimenetelmillä.

Ensiksi kehitetään vastakkainasettelullisen riskianalyysin (adversarial risk analysis, ARA) menetelmiä, joiden avulla tavoitteellisten omaa etuaan ajavien vastustajien päätösprosesseja voidaan mallintaa. Kehitetyt menetelmät mahdollistavat todennäköisyyksiin pohjaavan analyysin myös silloin, kun vastustajan käytettävissä olevaa tietoa ja hänen tavoitteitaan ei tunneta hyvin. Tästä on hyötyä tilanteissa, joissa vastustajan tavat analysoida ongelmaa ovat vaikeita arvioida ja hänen tavoitteensa jopa tarkoituksellisesti hämärän peitossa, mikä on yleistä esimerkiksi sotilastoimintaa analysoitaessa. Väitöskirjassa myös näytetään, miten taistelumallinnukseen kehitetyt mallinnus- ja simulaatiotyökalut voivat toimia osana vastakkainasettelullista riskianalyysiä. Näin olemassa olevilla työkaluilla voidaan vastata uuden tyyppisiin kysymyksiin, kuten kuinka omat päätökset voivat vaikuttaa vastustajan toimintaan tai kuinka yksittäisten joukkojen kohtaamiset voivat vaikuttaa taistelun kulkuun laajemmin.

Toiseksi kehitetään uusi ristivaikutusanalyysimalli (cross-impact analysis, CIA) tulevaisuuden epävarmuuksien jäsentelyyn ja ennakointiin asiantuntija-arvioiden perusteella. Näiden ristivaikutusarvioiden kanssa käytettäväksi on luotu kaksi laskennallista lähestymistapaa. Yksi käyttää arvioita todennäköisyyksien ja ristivaikutuksien ylä- ja alarajoista muodostamaan arvion koko systeemin riskistä tai odotusarvoisesta hyödystä. Toinen laskee saatavilla olevien todennäköisyys- ja ristivaikutusarvioiden perusteella yhteisjakauman eri epävarmuustekijöiden tapahtumien yhdistelmille sekä epävarmuustekijöitä kuvaavan Bayes-verkon. Tällaiset asiantuntija-arvioihin pohjaavat lähestymistavat ovat hyödyllisiä silloin, kun tilastollista tai simulointiaineistoa ei ole saatavilla. Näin voi käydä esimerkiksi silloin, kun tuloksia tarvitaan nopeasti tai kun epävarmuuksia on vaikea arvioida, koska ne liittyvät verrattain etäiseen tulevaisuuteen tai ihmisten toimintaan.

Avainsanat riskianalyysi, peliteoria, päätösanalyysi, skenaarioanalyysi, ristivaikutusanalyysi

ISBN (painettu) 978-952-64-1260-3 **ISBN (pdf)** 978-952-64-1261-0

ISSN (painettu) 1799-4934 **ISSN (pdf)** 1799-4942

Julkaisupaikka Helsinki **Painopaikka** Helsinki **Vuosi** 2023

Sivumäärä 122 **urn** http://urn.fi/URN:ISBN:978-952-64-1261-0

Preface

Working on this dissertation has been a great adventure, and one that I would have never embarked upon if it had not been for my supervisor Ahti Salo. Without him, this dissertation would almost certainly have never been written, and I feel that he deserves my utmost gratitude. The journey has not always been easy, but Ahti for his part has always made sure his students have had all the tools necessary for success.

I feel like I should also thank Esa Lappi and Bernt Åkesson. If they had not given me the first push all those years ago, I am not sure I would have ever ended up doing work related to military operations research. Had that not happened, this dissertation would most likely have ended up looking quite different. Of course, that was only the start of a very long path, and along the way, I have met so many wonderful collaborators, coworkers, and students. They have made my work so much more enjoyable and interesting, and I extend my gratitude to all of them.

Finally, I want to thank my mother Outi, my father Terho, and my grandparents Sirkka and Lasse. I have always been able to count on your support, and that has been invaluable.

Espoo, April 26, 2023,

Juho Roponen

Contents

List of Publications

This thesis consists of an overview and of the following publications which are referred to in the text by their Roman numerals.

I Juho Roponen and Ahti Salo. Adversarial risk analysis for enhancing combat simulation models. *Journal of Military Studies*, 6(2), 82-103, December 2015.

II Juho Roponen, David Ríos Insua, and Ahti Salo. Adversarial risk analysis under partial information. *European Journal of Operational Research*, 287(1), 306-316, November 2020.

III Ahti Salo, Edoardo Tosoni, Juho Roponen, and Derek W. Bunn. Using cross-impact analysis for probabilistic risk assessment. *Futures & Foresight Science*, 4(2), e2103, September 2021.

IV Juho Roponen and Ahti Salo. A probabilistic cross-impact methodology for explorative scenario analysis. Submitted to *Futures & Foresight Science*, December 2022.

Author's Contribution

Publication I: "Adversarial risk analysis for enhancing combat simulation models"

Roponen is the primary author. He wrote most of the manuscript and implemented the computational model for the case example. Salo helped finalize the manuscript.

Publication II: "Adversarial risk analysis under partial information"

Roponen is the primary author. He designed and implemented the methods and wrote most of the manuscript. Ríos Insua helped with structuring and writing the methodology. Salo helped finalize the manuscript.

Publication III: "Using cross-impact analysis for probabilistic risk assessment"

Salo is the primary author. Roponen helped write the section on the relationship between cross-impact statements and scenario probabilities, created the optimization model for the case study, and calculated the results. Tosoni provided the data and wrote the section on the case study. Bunn helped finalize and revise the manuscript.

Publication IV: "A probabilistic cross-impact methodology for explorative scenario analysis"

Roponen is the primary author. He designed and implemented the methods, collected the data and carried out the computations for the case study, and wrote most of the manuscript. Salo helped finalize the manuscript.

List of Figures

1. Introduction

1.1 Interpretation of probability

Humankind has always been surrounded by uncertainties. Yet, quantifying that uncertainty with probabilities is a relatively recent innovation (Bernstein, 1996). Still, people have long thought about uncertainty, for already ancient Roman and Jewish laws of evidence include degrees of proof and presumptions to deal with uncertainty in the court of law (Franklin, 2001). Yet, the first written record of probability calculus comes from the 1560s, when Gerolamo Cardano wrote about the sum totals when rolling three dice and how the odds of those totals arose from the combinations of dice rolls that could produce them (Bellhouse, 2005).

All the earliest examples of probability calculus focus on games of chance, including writings of Cardano, Pierre de Fermat, Blaise Pascal, Christiaan Hyugens, and even Galileo (Hacking, 2006). Because fair dice or decks of cards provide a discrete set of possible outcomes which can all be assumed equally likely, the underlying probabilities are easily understood. More challenging mathematical problems arise when the number of cards or dice increase, but this does not require substantial changes in problem framing. However, as events with equally likely outcomes rarely appear outside games of chance these early advances in probability theory found little practical use (Bernstein, 1996).

Probabilities cannot be observed or measured directly, which may in part explain why it took so long for probability theory to rise to prominence. Advances in the field of statistics eventually led to new applications for probability theory, finding uses in insurance pricing and policy decisions (Bernstein, 1996). Statistics and probability seemed like such a perfect match, that for a time, probability was widely interpreted as the frequency of a specific outcome when a trial was repeated infinitely many times. Whilst this frequentist interpretation explains well the probabilities of dice rolls and variations found in statistics, it is not particularly helpful for determining the probabilities of future events (French, 1986). Modern simulation models have made it possible to conduct statistical analysis about the future in well-understood physical systems by varying the initial conditions or simulation parameters (Hammersley, 2013)—an approach that is used in meteorological forecasting (Wilks and Wilby, 1999) and estimating the effects of weapon systems (Brandstein and Horne, 1998), for example. Yet, not all systems are easy to simulate, least of all human behavior.

Because of the difficulty or sheer impossibility of estimating the frequency or propensity of future events (or human activity) in many contexts, probabilities can instead be treated as degrees of beliefs in an event (Corfield and Williamson, 2001; Howard and Abbas, 2016; French, 1986; Raiffa, 1968). This interpretation is called Bayesian probability in honor of 18th-century mathematician Thomas Bayes, who first presented a method for updating beliefs about probabilities that is now called Bayesian inference (Bernstein, 1996). The subjective probability interpretation itself, however, should perhaps be more accurately attributed to Pierre-Simon Laplace (2012). Nowadays, subjective Bayesian probabilities have become the norm in fields like game theory and decision analysis, which deal with human decision-making in particular (Corfield and Williamson, 2001).

It is a matter of philosophical debate, whether the true nature of probability is statistical, subjective, or simply some hidden physical property (Kyburg and Smokler, 1980). Still, all these interpretations agree that probabilities can be treated in a rigorous mathematical manner following the rules of probability theory, a branch of mathematics with a well-defined set of axioms governing probabilities (Kallenberg, 1997). Thus, the exact probability interpretation rarely affects the validity of mathematical methods but can have implications on how the probabilities should be assessed and interpreted (Kyburg Jr, 2012). The author of this dissertation subscribes to the subjective Bayesian school of thinking.

1.2 Decision theory

The primary reason that probability theory is so widely applicable is that humans live in a world full of uncertainties. This is not an artifact of the modern world but has existed throughout history (Bernstein, 1996). Humans even seem to have evolved to have some innate understanding of probabilities (Fontanari et al., 2014). While the specific problems faced by humans have changed with the transition from hunter-gatherers to modern societies, uncertainty has not disappeared.

The formal answer guiding choices between uncertain alternatives has been known "ever since mathematicians first began to study the measurement of risk" (Bernoulli, 2011). Laplace (2012) called it mathematical hope.

Nowadays it is known, among other names, as mathematical expectation or expected value. In its simplest form, it is the sum of possible outcomes x_i weighted by their probabilities p_i

$$\mathrm{E}[X] = x_1 p_1 + x_2 p_2 + \ldots + x_n p_n, \tag{1.1}$$

but extensions also exist for countably and uncountably infinite sets of possible outcomes (Kallenberg, 1997).

Maximizing expected profits is at the core of most of risk analysis, insurance mathematics, and business optimization. However, for the maximization of mathematical hope to become a truly universal answer to decision problems under uncertainty, one more innovation was needed. In their book Theory of Games and Economic Behavior (von Neumann and Morgenstern, 1947), John von Neumann and Oskar Morgenstern laid the foundation for the field of decision theory (as well as game theory). Introducing what is today known as von Neumann–Morgenstern (VNM) utility theorem they postulated a set of axioms describing a rational decision-maker. When faced with risky outcomes this rational decision-maker should choose an alternative that maximizes the expected utility, defined as

$$\mathrm{E}[U] = u(x_1)\, p_1 + u(x_2)\, p_2 + \ldots + u(x_n)\, p_n, \tag{1.2}$$

where the utility function u provides a measure for the decision-maker's preferences over the possible outcomes x_i.

The VNM utility theory has been the target of criticism since its inception, sometimes even quite unjustly (Fishburn, 1989). Still, the theory has been very influential at the heart of the field of decision theory (Peterson, 2017), and even some of the critics are advocating for improving or evolving the theory instead of abolishing it (Kahneman and Tversky, 1979; Schoemaker, 1982; Fishburn, 2013). There is a wide range of literature on the study of utility functions and their role in human decision-making, but will not be discussed further here, because this dissertation's primary focus is on the other part of equation (1.2), that is, the probability distribution.

1.3 Objectives of the dissertation

This dissertation develops mathematical methods for characterizing probability information about uncertainties in order to support decision-making. Publications I and II focus primarily on the analysis of uncertainties arising from the competing activity of other decision-makers, whereas publications III and IV focus on quantifying the uncertainties based on expert judgments.

Publication I explores how adversarial risk analysis (ARA) could be used to model strategic and tactical decision-making in military combat modeling, where uncertainty arising from the decision processes of different actors has traditionally not been incorporated in the analysis. Typically, these decisions are modeled based on either very simplified game theoretical equilibriums or just expert opinion, leaving little to no room for uncertainty.

Publication II presents new computational methods for performing ARA without assigning probability distributions for all the involved uncertainties. This is useful especially in security contexts because it avoids assigning probability distributions to adversaries' actions and utilities. Because it is difficult to estimate the utility function of even a cooperative decision-maker, it can be almost impossible to accurately estimate the utilities of an adversarial decision-maker.

Publication III develops a new approach to cross-impact analysis by mapping expert judgments into corresponding probability bounds to different system outcomes. These probabilities are then used to establish upper and lower bounds for the system risk and other performance indicators. This approach makes it possible to form conservative estimates about system safety even if precise information about associated probabilities is not readily available.

Publication IV presents a computational approach to using CIA expert judgments, which may be imprecise and contradictory, to establish a probability distribution for possible system outcomes. This can facilitate probabilistic analysis based on future events and other difficult-to-model systems, like the ones often found in ARA.

2. Methodological foundations

The language used to describe decision problems is continuously evolving and quite diverse (Keith and Ahner, 2021), so it is necessary to first explain some of the terminology used throughout this chapter. *Probability* is taken as representing a degree of belief in an event, and an *event* can be any statement about the state of reality, for example, "Tomorrow it will rain." or "Galileo died in December.". A *random variable*[1] is a division of reality into multiple possibilities. These possibilities, called (the random variable's) *outcomes*, are mutually exclusive and jointly exhaustive events, so exactly one of them is guaranteed to always happen. A random variable could be for example "Tomorrow's highest temperature" with outcomes {"$<0°$C", "$\geq 0°$C and $\leq 20°$C", "$>20°$C"}. The random variables included in the analysis should always be chosen to be useful for characterizing the decision problem (Howard and Abbas, 2016).

2.1 Probability theory

The definitions here broadly follow (French, 1986) and (Kallenberg, 1997), although some of the terms and notation used are different. To start out, let $(\Omega,\mathscr{F},P)$ be a probability space. For the sake of simplicity, we assume that the sample space Ω is countable i.e. finite or countably infinite. Because Ω is countable, all of its subsets can be included in the event set, and thus the event set $\mathscr{F} = 2^{\Omega}$. The set of outcomes $\mathscr{X}$ for random variable X is defined as a countable partition of $\Omega = \bigcup_{x\in\mathscr{X}} x$. This means that $\mathscr{X} \subseteq \mathscr{F}$, and that the outcomes $x \in \mathscr{X}$ are mutually exclusive and collectively exhaustive. Because $\mathscr{X} \subseteq \mathscr{F}$, $P(x)$ is defined for all $x \in \mathscr{X}$, and also for unions and

[1]Various names exist for the same concept in the literature, for example, random variable (Kallenberg, 1997), random event (Harsanyi, 1967), key factor (Bunn and Salo, 1993), lottery (Raiffa, 1968; Myerson, 1997), distinction (Howard and Abbas, 2016), and uncertainty factor (Seeve and Vilkkumaa, 2022). Ultimately, the term random variable is used here to keep the terminology as familiar to most readers as possible. Notably, Publications III and IV primarily use the term uncertainty factor instead.

intersections of outcomes.

Because the outcome sets are countable, the random variables are discrete and defined as a function $X : \Omega \to \mathscr{X}$ such that the preimages $\mathscr{X}^{-1}(x) = x$, for all $x \in \mathscr{X}$. Therefore, random variables describe the possible states of the world by dividing them into specific outcomes. Usually, the focus of the analysis is on the induced distribution $P \circ X^{-1}$, and the choice of Ω plays little to no role. (Kallenberg, 1997).

It is convenient to also define the sets of possible decision outcomes $\mathscr{D}$ as countable partitions of $\Omega = \bigcup_{d \in \mathscr{D}} d$. Informally, this means that making decision d implies that we then live in a world we event d happens. This ensures that for example conditional probabilities such as

$$p_D(\omega|d) = \frac{p_D(\omega \cap d)}{p_D(d)}. \tag{2.1}$$

are defined when x is an outcome of a random variable and d is a decision alternative. We use P_D to denote the beliefs specific to decision-maker D when the distinction is necessary.

From the perspective of decision-maker D, there is normally no uncertainty about the outcome, so $P_D(d) = 1$ if they choose the decision alternative d. However, $P_A(d)$ may not be certain from the perspective of another decision-maker A, if they are unable to directly observe the decision. Defining the decision alternatives as events in the probability space means that all decision-makers use the same Ω and $\mathscr{F}$ and only different P. It also provides an easy way to define mixed decision strategies if necessary.

2.2 Decision theory

The basic decision problem examined in this dissertation is the expected utility maximization for a rational decision-maker D

$$\max_{d \in \mathscr{D}} \sum_{\omega \in \Omega} p_D(\omega|d)\ u_D(\omega), \tag{2.2}$$

where decision set $\mathscr{D}$ contains all of D's decision alternatives. D's probability estimate of ω given decision d is denoted with $p_D(\omega|d)$. D's utility is represented with a VNM utility function $u_D : \Omega \to \mathbb{R}$.

Real-life problems, however, are not always modeled with just a single conditional probability distribution, but often involve multiple decisions and other random variables (Raiffa, 1968). The utility function can also be expressed as a function of random variable outcomes instead of sample space Ω to better tie it to the problem structure (French, 1986), resulting in a utility function of form

$$u(\omega) = f\left(X_1(\omega),\ X_2(\omega), \ldots,\ X_N(\omega)\right). \tag{2.3}$$

In fact, this is often preferable in practice because it is prohibitively difficult to establish a utility function over a very large Ω otherwise.

2.3 Games

Decision problems with multiple decision-makers whose interests do not align are often modeled as games. In this context, a game is a collection of rules that describes the decision-makers' decision alternatives, available information, and random variables (Myerson, 1997). Incorporating another decision-maker changes the decision-maker D's expected utility to

$$\mathrm{E}[U_D] = \sum_{a \in \mathscr{A}} \sum_{\omega \in \Omega} p_D(\omega|a,d)\ p_D(a|d)\ u_D(\omega), \tag{2.4}$$

where $a \in \mathscr{A}$ represents the decision made by the other decision-maker, henceforth referred to as Adversary or A.

Finding the best decision alternative d now requires determining how D believes A will react to the changing environment as represented by $p_D(a|d)$. Without detailing the game, very little can be said about $p_D(a|d)$, because it depends on the information the decision-makers act on. In adversarial risk analysis, the problem is solved by assigning probabilities to the possible Adversary types represented by pair $T_A = (u_A, p_A)$, corresponding to A's utility function and beliefs about probabilities respectively (Banks et al., 2015).

Whilst it would be technically possible to include the Adversary's type in the same probability space as the decision problem, for the sake of simplicity its probability space will be denoted here with $(\mathscr{T}_A, F^T, P_D^T)$, where the sample space $\mathscr{T}_A$ is the set of possible Adversary types T_A. The probabilities are denoted with P_D^T to emphasize that these are D's subjective beliefs about the Adversary type.

Assuming that the Adversary is also a rational decision-maker, their decision can now be determined for each possible Adversary type

$$a(T_A) = \operatorname*{argmax}_{a \in \mathscr{A}} \sum_{\omega \in \Omega} p_A(\omega|a)\ u_A(\omega). \tag{2.5}$$

Here, it is assumed that the Adversary's decision $a(T_A)$ has a unique solution, but that is not always true. Analyzing the Adversary's decision problem may produce one or multiple optimal decisions a, or there may not be an optimum at all if set $\mathscr{A}$ is not finite.

The original decision problem can now be solved as

$$\max_{d \in \mathscr{D}} \sum_{T_A \in \mathscr{T}_A} p_D^T(T_A) \sum_{\omega \in \Omega} p_D(\omega|a(T_A),d)\ u_D(\omega), \tag{2.6}$$

assuming that mixed (randomized) decision strategies are disallowed. Incorporating mixed strategies would change the decision alternatives of all decision-makers into probability distributions over specific actions d and a, but otherwise, the problem would remain similar.

In practice however, it is difficult to assign well-founded probability distributions for $p_D^T(T_A)$, $p_D(\omega|a(T_A),d)$ and especially $p_A(\omega|a)$, because a rational Adversary should be expected to perform an analysis of their own, effectively mirroring what is done in (2.5) and (2.6). This translates the problem into a Bayesian game (Harsanyi, 1967), which cannot be solved without exploring the information upon which the decisions are based on. Thus, complex games are often not studied purely algebraically, but also incorporate graphical models that show what information is available at each stage of the decision process (Myerson, 1997).

2.4 Graphical models

Visual models are an integral part of analyzing complex decision problems. They provide an easy-to-read representation of the information and interdependence structure of the decision problem and are easier to construct and interpret than purely algebraic representations. Some of the simplest and most widely used visual representations are decision trees (see, for example French, 1986; Raiffa, 1968), and in the case of multiple decision-maker systems, game trees (Myerson, 1997), but they grow in size exponentially as the problem complexity increases. Here we opt to use graph-based models instead that provide largely the same information as decision trees but do not grow impractically large as quickly.

These networks use directed acyclic graphs to represent dependencies between random variables. The random variables are often chosen corresponding to some physical systems or easily observable system outputs. In more abstract problems, like those that concern warfare, counter-terrorism, or foresight, the random variables will be less concrete, but they are still chosen in a way that supports analyzing the decision problem at hand (Howard and Abbas, 2016).

2.4.1 Bayesian networks

A Bayesian network (Pearl, 1985) consists of a graph $G = (\mathcal{V}, E)$ that is a pair consisting of nodes (vertices) $\mathcal{V}$ that correspond to random variables and edges E that describe the conditional dependencies between the variables. With a slight abuse of notation, we will use X_i to denote both the random variables and the associated nodes of the network, so we write $\mathcal{V} = \{X_1, X_2, ..., X_N\}$. The set of edges E consists of ordered pairs of distinct nodes $E \subseteq \{(X,Y) | (X,Y) \in \mathcal{V}^2 \text{ and } X \neq Y\}$. Because we only discuss directed (simple) graphs, any mentions of edges refer to directed edges, and the existence of $(X,Y) \in E$ implies $(Y,X) \notin E$.

Figure 2.1 shows a simple Bayesian network. The circles represent nodes and the connecting arrows represent edges. The edges indicate probabilistic dependencies between the random variables, i.e. an edge from node X_1 to node X_2 implies that the conditional probability $P(X_2|X_1)$

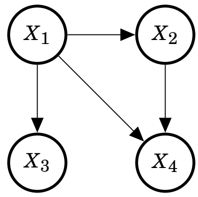

Figure 2.1. A Bayesian network.

differs in some way from probability $P(X_2)$ for some events in $\mathscr{X}_1$ and $\mathscr{X}_2$. Conversely, the lack of a connecting edge implies conditional independence, i.e. the conditional probability distributions of the two random variables conditioned on each of their respective incoming edges are independent. The conditional independence relation is context-specific and depends on which random variables are included in the network and also on the direction of the edges.

The probability information associated with a Bayesian network is expressed as conditional probabilities. The probability of every outcome is conditioned on other random variables connected by an incoming edge. For example, in the Bayesian network from Figure 2.1 random variable X_2 would have its probability distribution encoded as $P(X_2|X_1)$. The outcomes of X_3 are not included, because it does not share an edge with X_2, and neither are the outcomes of X_4, because the edge between the two nodes is directed from X_2 to X_4.

Following from the law of total probability and the definition of conditional independence, these conditional probabilities are sufficient to specify the probability of any combination of outcomes $(x_1, x_2, ..., x_N)$. Using the example from Figure 2.1 again, we get

$$P(x_1, x_2, x_3, x_4) = P(x_1)P(x_2|x_1)P(x_3|x_1, x_2)P(x_4|x_1, x_2, x_3) \tag{2.7}$$

$$= P(x_1)P(x_2|x_1)P(x_3|x_1)P(x_4|x_1, x_2). \tag{2.8}$$

From the computational perspective, these conditional probabilities are also convenient because they require storing far less information than probabilities of outcome combinations separately, assuming the network is sparse enough.

2.4.2 Influence diagrams

Influence diagrams (Howard and Matheson, 2005) can be treated as an extension of Bayesian networks despite predating them conceptually (Howard and Matheson, 1984; Verma and Pearl, 1988; Pearl, 2005). Like a Bayesian network, an influence diagram also consists of a graph $G = (\mathcal{V}, E)$, but unlike in a Bayesian network some nodes of the graph represent decisions $\mathcal{V}_D \subset \mathcal{V}$ and decision-maker utility $\mathcal{V}_U \subset \mathcal{V}$. In other words, an influence diagram with neither decision nor utility nodes is just a Bayesian network

(Kjaerulff and Madsen, 2008).

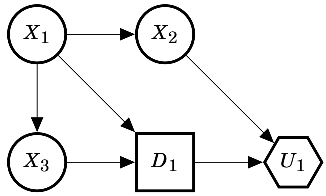

Figure 2.2. An influence diagram.

Figure 2.2 depicts an example influence diagram. The circles represent random variables as they do in Bayesian networks, the squares represent decisions, and the hexagons represent utility to the decision-maker. The incoming edges to decisions indicate that the outcomes of these connected nodes are known at the time of the decision, whilst incoming edges to utility nodes indicate which decision and random variable outcomes are used to calculate the utility of the decision-maker.

The graphical representation is accompanied by a description detailing if the incoming edges affect the decisions beyond providing information, for example, if they limit the available decision alternatives in some way. Typically, with influence diagrams, it is assumed that all the information the decision-maker had access to during earlier decisions, as well as the decision outcomes, are known when making later decisions (Shachter, 1986; Tatman and Shachter, 1990), but this assumption can be expressly omitted in some cases (Mauá et al., 2012; Kjaerulff and Madsen, 2008).

Unlike random variables and decisions, utility nodes do not involve any uncertainty. They are defined as VNM-utility functions over random variables and decisions that are connected to the utility node. Typically, an influence diagram will have exactly one utility node, but sometimes multiple are used to represent separable components of the utility function (Tatman and Shachter, 1990).

2.4.3 Multi-agent influence diagrams

Multi-agent influence diagrams (MAIDs) are influence diagrams that can be used to represent games by including decisions, uncertainties, and utilities of multiple decision-makers in a single graph (Koller and Milch, 2003). This provides a more compact visual representation of a game than a game tree would, whilst still making it possible to denote the order of decisions and uncertain events as well as the flow of information.

In the MAID in Figure 2.3 colors are used to differentiate between agents. Here, decision-maker D's decisions as well as utilities and uncertainties only relevant to them are colored white, whereas vertices associated with Adversary A are colored gray. Random variable X_2 affects the utility of both D and A, so it is colored with white and gray stripes. Random variable

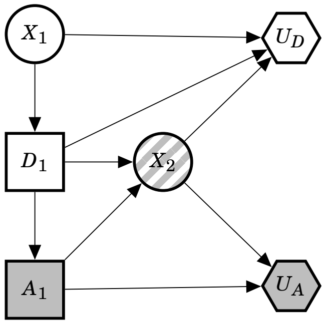

Figure 2.3. A multi-agent influence diagram.

X_1 on the other hand is entirely irrelevant to A's decision-making, so it is colored white.

Sometimes when using MAIDs, coloring the random variables is omitted because the same relevant information can be deduced from the network structure and the accompanying descriptions of uncertainties, decisions, and utilities. However, coloring all nodes helps separate the decision problems of different actors. Figure 2.4 shows the decision problems of the two agents separately. Producing these influence diagrams is very straightforward when the original MAID is colored (Ortega et al., 2019). The other agent's decisions are replaced with random variables, and the random variables and the utilities associated only with other agents are removed.

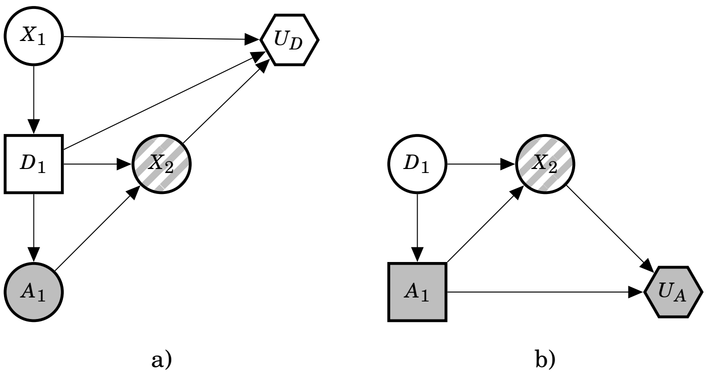

Figure 2.4. The decision-maker's problem a) and the adversary's problem b) as separate influence diagrams.

2.5 Cross-impact analysis

Cross-impact analysis (CIA) encompasses several methods built on the ideas presented in the seminal work of Theodore Gordon and Olaf Helmer in the 1960s (Gordon, 1994). CIA methods are crafted for the purpose of examining and characterizing interdependencies that exist between random variables, also referred to as uncertainty factors. By having experts rate these dependencies' magnitudes and directions on a numerical scale, it is then possible to draw conclusions about their joint probabilities.

Whilst there are almost as many ways of measuring cross-impacts as there are cross-impact methods (see e.g. Alter, 1979; Amer et al., 2013; Bishop et al., 2007), this dissertation exclusively uses the definition first presented in Publication III, referred to as the cross-impact multiplier definition. Cross-impact multiplier for outcomes x_1 and x_2 of random variables X_1 and X_2 respectively is defined as

$$C_{x_1x_2} = \frac{P(x_1 \cap x_2)}{P(x_1)P(x_2)}. \tag{2.9}$$

It is called the cross-impact multiplier, because it describes the relative change in probability of outcome x_1 when x_2 is known to happen compared to when nothing is known about x_2. This is because

$$C_{x_1x_2} = \frac{P(x_1 \cap x_2)}{P(x_1)P(x_2)} = \frac{P(x_1|x_2)}{P(x_1)}. \tag{2.10}$$

Compared to other cross-impact approaches, the cross-impact multipliers have some distinct advantages.

i They facilitate estimating numerical probabilities and are thus compatible with risk and decision analysis methods as well as graphical probability models.

ii They are symmetric by design, i.e. $C_{x_1x_2} = C_{x_2x_1}$ as seen from (2.9), which means that dependencies do not need to be evaluated twice for every outcome pair.

iii They avoid interacting directly with conditional probabilities, which can be difficult for non-experts to estimate (Pollatsek et al., 1987).

3. Research Contributions

This dissertation presents new methods for supporting decision-making under uncertainty, especially in problems relating to safety and security. Publications I-II focus on using adversarial risk analysis (ARA) to model the decision processes of adversarial actors. Publications III-IV develop probabilistic cross-impact methods to support risk evaluation and probability estimation based on expert judgments. The contributions are summarized in Table 3.1.

3.1 Publication I

ARA combines methods of statistical risk analysis and game theory to help evaluate risks and choose countermeasures against threats posed by intelligent and potentially malignant actors. Many of the earliest ARA applications have focused on counterterrorism. However, despite the long history of game theory and computational models in the military, ARA has not been widely applied in combat modeling (at least publicly). Publication I identifies ways to combine ARA methods with existing combat modeling tools to broaden the scope of analyses that can be performed. Specifically, Publication I discusses how ARA could serve to combine results from low-level simulations to form a picture of how the success of individual units could affect the wider conflict. The publication also offers a simple example of how ARA can be used to evaluate the value of information and the importance of operational secrecy.

3.2 Publication II

Evaluating the rationale of other decision-makers poses a persistent challenge in applying ARA and other methods based on game theory to real-world problems. Evaluating the utility function of a cooperative party is challenging, but finding reliable information about the beliefs and pref-

Table 3.1. Summary of publications.

Publication	Objectives	Methodology	Results
I	To explore and demonstrate the potential of ARA in military applications.	Adversarial Risk Analysis, Combat Modeling, Simulation	ARA can be combined with existing combat models to analyze questions that would be outside the original model's scope, including decision problems, impacts on the conflict on a larger scale, and the value of hidden information.
II	To develop new ARA methods to analyze problems in which probabilities about the adversary or some other aspect of the system are not known.	Adversarial Risk Analysis, Stochastic Dominance, Combat Modeling, Simulation	The developed method enables solving all ARA problems represented by regular influence diagrams using stochastic dominance and decision rules when exact utility functions or probability distributions are not available. A case study demonstrates the use of the developed method for military planning.
III	To provide a cross-impact interpretation founded on probability theory for use with risk analysis.	Cross-impact Analysis, Risk Analysis, Quadratic Programming	A new definition for probabilistic cross-impacts founded on probability theory is presented. Applicability to risk analysis is demonstrated with an optimization method and a case study.
IV	To develop a method for computing scenario probability distribution based on cross-impact information.	Cross-impact analysis, Scenario Analysis, Least Squares Approximation	A new optimization method, which utilizes cross-impact information to compute joint probability distributions for random variables. A Bayesian network can also be constructed based on the computed probabilities. Applicability to real-world problems is demonstrated with a case study.

erences of adversaries can be near impossible. To avoid having to make unrealistically precise predictions about adversaries' thought processes, Publication II develops methods for characterizing their likely actions based on more general assumptions. Publication II shows how the concepts of partial information, stochastic dominance, and decision rules can be used instead of some or all of the probability distributions and utility functions to analyze adversarial risks. The contributions are demonstrated with a realistic case study about choosing and deploying countermeasures to unmanned aerial vehicles.

3.3 Publication III

Risk analysis of complex systems calls for the identification, characterization, and analysis of numerous possible future events and developments that may negatively impact the system. The task is further complicated by the fact that these uncertainties can also depend on one another. However, looking individually at every possible scenario that can be formed as a combination of these outcomes quickly becomes infeasible when the number of random variables increases. Cross-impact analysis offers a tool for estimating how the perceived probabilities of random variables change based on the outcomes of others. Publication III offers a cross-impacts definition that is founded on probability theory and admits several kinds of probabilistic statements about dependencies between the uncertainty factors. The publication also describes how the statements can be transformed into optimization constraints and used to calculate upper and lower bounds for the overall risk level of the system. The approach is illustrated with an example case about the risk analysis of nuclear waste repositories.

3.4 Publication IV

Estimating the probability distributions of the different random variables is one of the main challenges in producing probabilistic forecasts. Simulation models, such as the ones used in weather forecasting and military combat modeling, can be used to quantify future uncertainties governed by chance. However, modeling uncertainty stemming for example from human activity with simulations is often not feasible, and creating detailed simulation models is challenging and time-consuming. Thus, eliciting experts for their estimations about future uncertainties is often the only feasible approach. Still, eliciting information about systems with multiple interdependent random variables poses a challenge, because when the number of variables increases the number of possible interactions with them grows exponentially. Cross-impact methods manage this complexity by focusing only on the pairwise impacts between two random variables. Publication IV presents a new optimization approach that can synthesize these pairwise cross-impact statements to produce a joint probability distribution for the random variables. When combined with conditional independence information, the calculated probabilities can also be used to construct a Bayesian network to aid what-if analyses. The Publication also includes a case study focusing on the future of 3D-printing in military use.

4. Discussion

This dissertation develops new methods to account for uncertainties in support of decision-making in adversarial risk analysis (ARA) and probabilistic scenario analysis. Publications I and II focus on ARA methodology to quantify uncertainties caused by adversarial decision-makers with competing interests. Publications III and IV, on the other hand, present novel approaches for using cross-impact analysis (CIA) to quantify uncertainties associated with future events.

Despite having been developed relatively recently, ARA has already found numerous applications in counter-terrorism and cyber security. Still, much of the military combat modeling research does not use ARA or any other game-theoretic models for adversarial activity (Washburn et al., 2009), despite the great impact that adversaries' actions have on the effectiveness of tactics and weapon systems. As discussed in Publication I, ARA methods can be used to expand the range of analyses that can be performed using pre-existing combat modeling tools, incorporating small-scale encounters as a part of the bigger picture and evaluating the value of secrecy and information.

The shared common knowledge assumptions required by many game-theoretical models have been problematic in attempts to adapt game theory to combat modeling. Bayesian Nash equilibrium developed by Harsanyi (1967) provides the necessary tools for finding robust solutions for facing different types of adversaries, but coming up with a probability distribution over *adversary types* (representing their utility functions and available information) can be onerous. In Publication II, we show that even simple assumptions about adversaries (such as wanting to minimize casualties) can serve as a foundation for a game-theoretic analysis when interpreted as partial preference order relations over outcomes. Whilst this type of analysis cannot be used to predict the adversaries' actions precisely, some of their decision alternatives can be excluded as irrational. Thus, it is possible to find risk mitigation strategies that work against rational adversaries, even if the adversary's precise type or type's probability is not known.

Compared to ARA, CIA approaches quantifying uncertainty very differently. Cross-impact information elicited from experts describes how the likelihood of specific outcomes changes when an outcome of another random variable is known. There exist cross-impact methodologies that differ in almost everything except that basic idea (Alter, 1979; Amer et al., 2013; Bishop et al., 2007). Most of the earliest methods worked similarly to Monte Carlo simulation, drawing random events from the possible list of outcomes and then adjusting the probabilities of remaining outcomes based on the associated cross-impacts (e.g. Gordon, 1994; Dalkey, 1971; Helmer, 1981). More recently, several CIA methods have been developed that eschew probabilities entirely, and only measure the likelihood of outcomes appearing together in terms of how consistent their cross-impacts are (e.g. Weimer-Jehle, 2006; Seeve and Vilkkumaa, 2022). Whilst both of these approaches have their uses in exploring the future, developing scenarios, and fostering managerial thinking, they are not very compatible with either risk or decision analysis.

To facilitate risk analysis based on cross-impacts, Publication III presents a new cross-impact interpretation founded on probability theory. Called cross-impact multipliers, this new cross-impact interpretation, together with information about marginal probabilities of the associated random variables, can be used to determine upper and lower bounds for system risk. Thus it is useful, for example, in demonstrating compliance with regulatory risk bounds as well as in comparing different risk mitigation alternatives.

Whilst the primary focus in Publication III is on risk analysis, Publication IV takes the same cross-impact definition and presents methods for calculating a joint probability distribution for different scenarios formed as combinations of outcomes of random variables. It is also demonstrated, how the computed probabilities together with conditional independence information can be used to construct Bayesian networks, offering a useful tool for what-if type analyses.

This dissertation opens up new research directions as well. First, the presented methods could be tested with more empirical case studies. It would be interesting to try how compatible ARA and CIA are together. ARA is often the preferred method for modeling uncertainty from human activity, and CIA is good for estimating long-term technological and other developments. The two together could be used to analyze safety and security problems with long time horizons, such as investments into new weapon systems or the design of long-term nuclear waste repositories.

There is room for new methodological extensions as well. Expanding further on the methods presented in Publication II, it would be interesting to explore how the partial information approach could be expanded to also include non-sequential games, i.e. games with decisions whose outcomes are not observable before the next decision of the adversary. Although it is possible that the increased uncertainty would make it impossible to draw any useful conclusions about these games (Fishburn, 1978). Another potential research topic would be examining how different types of ambiguous preference models, such as the ones used by Danielson et al. (2014) or Salo and Punkka (2005), could be applied in ARA.

Expanding on the CIA methods presented in this dissertation, some work has already been done in using the computed probabilities or risks to choose scenarios for more detailed examination (Elfving, 2023). This way probabilistic and narrative scenario methods could be used together to combine some of the best aspects of both traditions. The analytical models help contextualize the scenarios and the narrative approaches provide depth and approachability. Selecting the right scenarios to focus on is also important in many modeling or simulation studies and offers another potential avenue for future research.

References

Steven Alter. The evaluation of generic cross-impact models. *Futures*, 11(2): 132–150, 1979.

Muhammad Amer, Tugrul U Daim, and Antonie Jetter. A review of scenario planning. *Futures*, 46:23–40, 2013.

David L Banks, Jesus M Rios Aliaga, and David Ríos Insua. *Adversarial Risk Analysis*. CRC, Boca Raton, FL, 2015.

David Bellhouse. Decoding Cardano's Liber de ludo aleae. *Historia Mathematica*, 32(2):180–202, 2005.

Daniel Bernoulli. Exposition of a new theory on the measurement of risk. In *The Kelly Capital Growth Investment Criterion: Theory and Practice*, pages 11–24. World Scientific, 2011.

Peter L Bernstein. *Against the Gods: The Remarkable Story of Risk*. John Wiley & Sons, New York, 1996.

Peter Bishop, Andy Hines, and Terry Collins. The current state of scenario development: an overview of techniques. *Foresight*, 9(1):5–25, 2007.

Alfred Brandstein and Gary Horne. Data farming: A meta-technique for research in the 21st century. *Maneuver Warfare Science*, 1998:93–99, 1998.

Derek W Bunn and Ahti A Salo. Forecasting with scenarios. *European Journal of Operational Research*, 68(3):291–303, 1993.

David Corfield and Jon Williamson. *Foundations of Bayesianism*, volume 24. Springer Science & Business Media, 2001.

Norman C Dalkey. An elementary cross-impact model. *Technological Forecasting and Social Change*, 3:341–351, 1971.

Mats Danielson, Love Ekenberg, Aron Larsson, and Mona Riabacke. Weighting under ambiguous preferences and imprecise differences in a cardinal rank ordering process. *International Journal of Computational Intelligence Systems*, 7(2):105–112, 2014.

Vikke Elfving. Scenario selection in probabilistic cross-impact analysis. Master's thesis, Aalto University, 2023.

Peter C Fishburn. Non-cooperative stochastic dominance games. *International Journal of Game Theory*, 7(1):51–61, 1978.

Peter C Fishburn. Retrospective on the utility theory of von Neumann and Morgenstern. *Journal of Risk and Uncertainty*, 2:127–157, 1989.

Peter C Fishburn. *The Foundations of Expected Utility*, volume 31. Springer Science & Business Media, 2013.

Laura Fontanari, Michel Gonzalez, Giorgio Vallortigara, and Vittorio Girotto. Probabilistic cognition in two indigenous mayan groups. *Proceedings of the National Academy of Sciences*, 111(48):17075–17080, 2014.

James Franklin. *The Science of Conjecture: Evidence and Probability before Pascal*. Johns Hopkins University Press, 2001.

Simon French. *Decision Theory: An Introduction to the Mathematics of Rationality*. Halsted Press, 1986.

Theodore J Gordon. *Cross-impact Method*, volume 4. American Council for the United Nations University, 1994.

Ian Hacking. *The Emergence of Probability: A Philosophical Study of Early Ideas About Probability, Induction and Statistical Inference*. Cambridge University Press, 2006.

John Hammersley. *Monte Carlo Methods*. Springer Science & Business Media, 2013.

John C Harsanyi. Games with incomplete information played by "Bayesian" players, I–III Part I. The basic model. *Management science*, 14(3):159–182, 1967.

Olaf Helmer. Reassessment of cross-impact analysis. *Futures*, 13(5):389–400, 1981.

Ronald A Howard and Ali E Abbas. *Foundations of Decision Analysis*. Pearson Education Limited, 2016.

Ronald A Howard and James E Matheson. *The Principles and Applications of Decision Analysis*. Strategic Decisions Group, Palo Alto, CA, 1984.

Ronald A Howard and James E Matheson. Influence diagrams. *Decision Analysis*, 2(3):127–143, 2005.

Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):363–391, 1979.

Olav Kallenberg. *Foundations of Modern Probability*. Springer, 1997.

Andrew J Keith and Darryl K Ahner. A survey of decision making and optimization under uncertainty. *Annals of Operations Research*, 300(2):319–353, 2021.

Uffe B Kjaerulff and Anders L Madsen. *Bayesian Networks and Influence Diagrams*. Springer, 2008.

Daphne Koller and Brian Milch. Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*, 45(1):181–221, 2003.

Henry E Kyburg and Howard E Smokler. *Studies in Subjective Probability*. New York, US: Robert E. Krieger Publishing Company, 1980.

Henry E Kyburg Jr. *The Logical Foundations of Statistical Inference*. Springer Science & Business Media, 2012.

Pierre-Simon Laplace. *A Philosophical Essay on Probabilities*. Courier Corporation, 2012.

Denis Deratani Mauá, Cassio P de Campos, and Marco Zaffalon. Solving limited memory influence diagrams. *Journal of Artificial Intelligence Research*, 44: 97–140, 2012.

Roger B Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1997.

Jorge Ortega, David Ríos Insua, and Javier Cano. Bi-agent influence diagrams from an adversarial risk analysis perspective. *European Journal of Operational Research*, 273(3):1085–1096, 2019.

Judea Pearl. Bayesian networks: A model of self-activated memory for evidential reasoning. In *Proceedings of the 7th Conference of the Cognitive Science Society, University of California, Irvine, CA, USA*, pages 15–17, 1985.

Judea Pearl. Influence diagrams—historical and personal perspectives. *Decision Analysis*, 2(4):232–234, 2005.

Martin Peterson. *An introduction to decision theory*. Cambridge University Press, 2017.

Alexander Pollatsek, Arnold D Well, Clifford Konold, Pamela Hardiman, and George Cobb. Understanding conditional probabilities. *Organizational Behavior and Human Decision Processes*, 40(2):255–269, 1987.

Howard Raiffa. *Decision Analysis: Introductory Lectures on Choices under Uncertainty.* Addison-Wesley, 1968.

Ahti Salo and Antti Punkka. Rank inclusion in criteria hierarchies. *European Journal of Operational Research*, 163(2):338–356, 2005.

Paul JH Schoemaker. The expected utility model: Its variants, purposes, evidence and limitations. *Journal of Economic Literature*, pages 529–563, 1982.

Teemu Seeve and Eeva Vilkkumaa. Identifying and visualizing a diverse set of plausible scenarios for strategic planning. *European Journal of Operational Research*, 298(2):596–610, 2022.

Ross D Shachter. Evaluating influence diagrams. *Operations Research*, 34(6): 871–882, 1986.

Joseph A Tatman and Ross D Shachter. Dynamic programming and influence diagrams. *IEEE Transactions on Systems, Man, and Cybernetics*, 20(2):365–379, 1990.

Thomas Verma and Judea Pearl. Influence diagrams and d-separation. Technical report, UCLA, Computer Science Department, 1988.

John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1947.

Alan R Washburn, Moshe Kress, et al. *Combat Modeling*, volume 139. Springer, New York, 2009.

Wolfgang Weimer-Jehle. Cross-impact balances: A system-theoretical approach to cross-impact analysis. *Technological Forecasting and Social Change*, 73(4): 334–361, 2006.

Daniel S Wilks and Robert L Wilby. The weather generation game: A review of stochastic weather models. *Progress in Physical Geography*, 23(3):329–357, 1999.

Publication I

Juho Roponen and Ahti Salo. Adversarial risk analysis for enhancing combat simulation models. *Journal of Military Studies*, 6(2), 82-103, December 2015.

Juho Roponen, M.Sc., Systems Analysis Laboratory, Dept. of Mathematics and Systems Analysis, Aalto University School of Science; juho.roponen@gmail.com
Ahti Salo, Professor, Systems Analysis Laboratory, Dept. of Mathematics and Systems Analysis, Aalto University School of Science, ahti.salo@aalto.fi

ADVERSARIAL RISK ANALYSIS FOR ENHANCING COMBAT SIMULATION MODELS

Abstract

Adversarial Risk Analysis (ARA) builds on statistical risk analysis and game theory to analyze decision situations involving two or more intelligent opponents who make decisions under uncertainty. During the past few years, the ARA approach–which is based on the explicit modelling of the decision making processes of a rational opponent–has been applied extensively in areas such as counterterrorism and corporate competition. In the context of military combat modelling, however, ARA has not been used systematically, even if there have been attempts to predict the opponent's decisions based on wargaming, application of game theoretic equilibria, and the use of expert judgements. Against this backdrop, we argue that combining ARA with military combat modelling holds promise for enhancing the capabilities of combat modelling tools. We identify ways of combining ARA with combat modelling and give an illustrative example of how ARA can provide insights into a problem where the defender needs to estimate the utility gained from hiding its troop movements from the attacker. Even if the ARA approach can be challenging to apply, it can be instructive in that relevant assumptions about the resources, expectations and goals that guide the adversary's decisions must be explicated.

Key words

Adversarial Risk Analysis (ARA), Combat Modeling, Simulation

Introduction

Adversarial risk analysis (ARA) combines statistical risk analysis and game theory to provide appropriate methods for analyzing decision making situations which involve two or more intelligent actors who make decisions with uncertain outcomes. Such situations are encountered, for example, in counter-terrorism and corporate competition (Rios Insua *et al.,* 2009).

Traditional statistical risk analysis was developed to assess and mitigate risks in contexts where the loss is governed by chance (or Nature), for instance in the management of complex technological systems like nuclear power plants and the design of insurance policies against natural disasters. Apart from risks caused by such chance events, ARA seeks to capture risks caused by the self-interested and possibly malicious actions of intelligent actors: thus, modelling the decision-making behavior of these actors is central to ARA. These kinds of decision models can be based, for example, on classical game theory (Myerson, 1991) or psychological considerations (Camerer, 2003).

Yet game theory is not an ideal tool for describing and predicting human behavior. Minmax solutions–in which each actor seeks to minimize his expected losses across all the actions that are available to his opponents–can lead to unrealistic solutions, because real opponents do not usually follow the minmax rationality principle. Minmax solutions are also often difficult to compute in real situations, and they necessitate strong assumptions about what common knowledge the actors share (Kadane & Larkey, 1982 and Meng *et al.,* 2014). Moreover, the solutions can be overly pessimistic, because the mitigation of the worst possible scenario (which may have an extremely low probability) will induce the actors to make choices that a human opponent would not realistically make.

ARA has many obvious uses in military organizations. Much of the recent ARA literature has focused on counterterrorism, and many of the proposed ARA approaches can be applied to support military decision making. Zhuang and Bier (2007), for example, apply game theory to devise strategies for allocating resources between the protection from an intentional attack, on one hand, and from natural disasters, on the other hand. ARA methods can also be used to guide the allocation of resources between strategically important targets as well as the investment planning of military equipment and projects. Uses of ARA in finance and procurement are relevant, too, because military organizations acquire products and services from external contractors.

In this paper, we do not survey the broad ARA literature in view of military applications. Rather, we discuss how ARA can be applied to enhance combat modeling or to complement it. Specifically, we examine how ARA can be used to model the effects of military deceit, and how it can be used to aggregate results from different simulations to model a longer chain of events.

The possibility of calculating the effects of military deception and its usefulness is one of the most promising ARA applications in combat modeling. Game theory has been applied to calculate the benefits of deceit before (Reese, 1980), but such applications are still rare. This is partly because the solutions of classical game theory presume that both sides have common knowledge about each other's goals and resources, which is not realistic when modeling deceit. ARA does not have this limitation. It can even be applied to calculate the usefulness of decoys and dummy systems, which makes it possible to estimate if these are worth the cost. Such estimation is very difficult if not impossible in most combat simulation models.

Using ARA to facilitate the simulation of longer chains of events holds promise, because simulation models are built to model combat on specific scale. For instance, simulation models which seek to accurately describe the combat between two tanks are ill-suited for modeling an entire battalion, and models for simulating fighting at the platoon level do not lend themselves well to the modeling an entire theater of operations. Still, with ARA it is possible combine results from several such simulation runs or even different simulation models to create a more encompassing optimization model. This can be very useful in stretching the limits of what can be done with existing simulation tools.

Modeling adversarial risks

In this section, we briefly describe how a situation in which there are adversaries whose actions affect each other's risks can be modeled. Our presentation builds largely on Rios Insua *et al.* (2009) who give a comprehensive presentation of ARA. For a good overview on how the ARA approach compares to classical game theory, we refer to Banks *et al.* (2011).

Risk analysis

The simplest form of a non-adversarial risk management problem is a situation in which the decision maker chooses one of the available decision alternatives whose costs are uncertain. This problem can be presented as an influence diagram as seen in Figure 1.

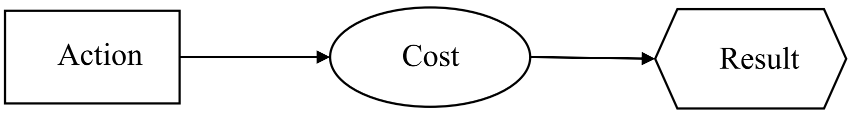

Figure 1: A simple influence diagram

An influence diagram is a directed acyclic graph with three kinds of nodes: rectangle shaped decision nodes, oval shaped uncertainty nodes, and hexagonal value nodes. Arrows pointing to value or uncertainty nodes indicate functional or probabilistic dependence, respectively. This means that the utility function at the value node depends on its immediately preceding nodes, and the probabilities associated with an uncertainty node depend are conditioned on the values of the immediately preceding nodes. Arrows which point to decision nodes indicate that the values of these preceding nodes are known at the time of the decision. (*cf.* Howard & Matheson, 2005)

Figure 1 shows a situation where the decision maker has to make a decision a from a set A of possible choices, represented by the rectangle. The cost c associated with this decision is uncertain and is modeled through the probability density function $\pi(c|a)$, represented by the oval node. The result is modeled by Von Neumann-Morgenstern utility function $u(c)$. The decision maker seeks to maximize the expected utility

$$\psi = \max_{a \in A}\left[\psi(a) = \int u(c)\pi(c|a)\, dc\right]. \qquad (1)$$

In practice, the costs of a particular action can be complex in that they can include both fixed and random terms. As a result, organizations seek to perform a risk assessment to better identify disruptive events, and to estimate their probabilities and associated costs.

Adversarial risks

We now consider a situation in which there are two adversaries (Attacker and Defender) whose decisions affect the risks that each faces. Figure 2 extends the influence diagram to include the adversary in a symmetrical situation in which the decisions of both parties affect the risks and costs that the other faces, and both seek to maximize their own expected utility. In this example, the roles are symmetric; but it is possible to model asymmetric scenarios as well by building asymmetric influence diagrams.

We denote the sets of possible actions of Attacker and Defender with A and D respectively. Their utility functions are $u_a(\cdot)$ and $u_d(\cdot)$. The sets containing their beliefs about different probabilities are P_a and P_d. As can be seen in the influence diagram in Figure 2, one of the nodes, Hazard, is common to both sides. It can represent uncertainties which affect parties, such as weather for example. The other cost nodes–which are not common–represent random costs for both parties which can be very different for the parties.

The expected utilities for both the Attacker and the Defender depend upon the actions of both. Specifically, by extending on (1), we obtain the Attacker's expected utility for choosing action $a \in A$ when the Defender chooses action $d \in D$

$$\psi_A(a,d) = \int u_A(c)\pi_A(c|a,d)\,dc, \qquad (2)$$

where $\pi_A(c|a,d)\epsilon P_A$ represents the Attacker's beliefs about his costs for the decision pair (a,d). It is noteworthy that these beliefs do not necessarily have to match reality, because we are only modeling the Attacker's decision. The expected utility for the Defender is analogous.

This representation of ARA matches normal form games in which both players take simultaneous decisions. One could also build an influence diagram that represents sequential games, such as Stackelberg games, in which the players make their moves alternately. The ARA methodology can be applied to solve such games, too (*cf.* Banks *et al.,* 2011 and Rios & Rios Insua, 2012).

Bayesian framework for ARA

A problem like the one presented in Figure 2 can be solved using classical game theory if the costs and utility functions of both players are common knowledge. However, if the players do not have correct and accurate information about the costs, resources, and goals of the adversary (which is often the case in reality), the Nash equilibrium solution does not exist.

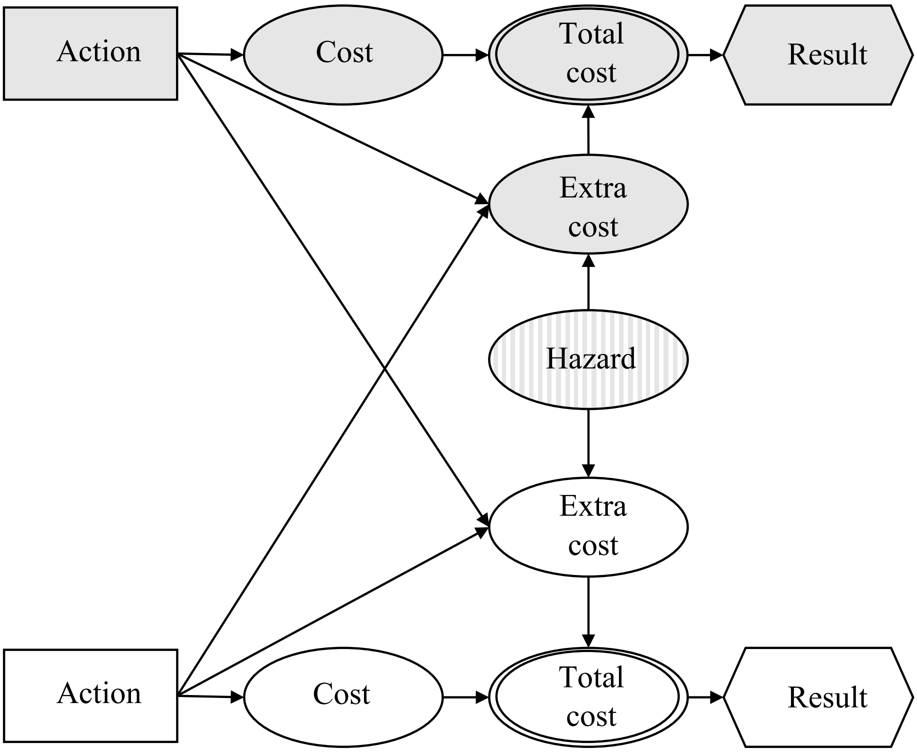

Figure 2: Influence diagram with an adversary

ARA solves this problem by using a Bayesian strategy to express uncertainty about the adversary's decision. From the Attacker's point of view, this means that the Defender's decision is a random variable as presented in Figure 3. To solve this problem, the Attacker needs more than just $\pi_A(c|a,d)\epsilon P_A$ and $u_A(c)$. Specifically, he also needs $p_A(d)$, which is the probability that the Defender chooses defense d as estimated by the Attacker. To find that, the Attacker is assumed to use mirroring to form an estimate of the Defender's utility function $u_D(c)$ and the Defender's costs $\pi_D(c|a,d)$. In other words, the Attacker assumes that the Defender acts rationally and that the Defender uses a similar approach to predict the actions of the Attacker.

If the Attacker tried to estimate the Defender's utility function and cost function by assuming that the Defender is doing exactly the same thing as what he is doing, the Attacker would need to think what the Defender thinks he thinks. To avoid infinite regress, the chain is usually cut here and the Attacker just forms an educated guess about the Defender's beliefs about the Attacker's estimated utilities and costs. In principle, this analysis could be taken even further, but usually this is not realistic.

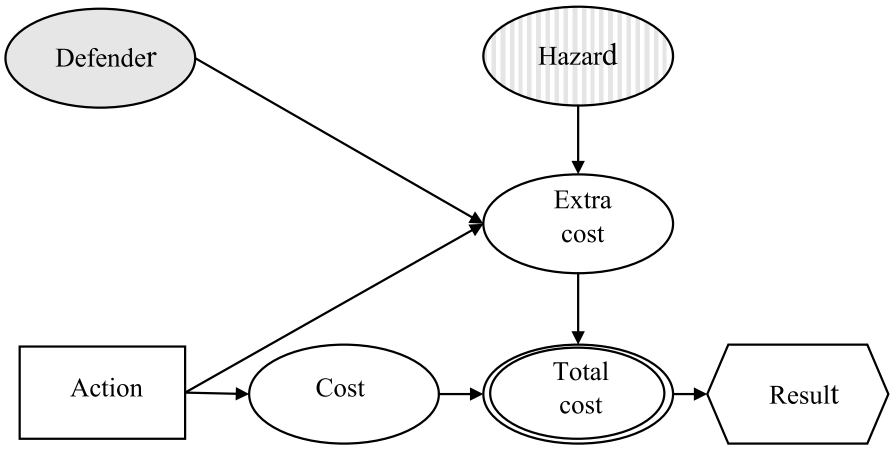

Figure 3: Influence diagram from the Attacker's point of view.

Alternative approaches for modeling adversary's decision making

The ARA methodology has analogues with Bayesian level-k thinking: specifically, our approach to modeling of the opponent's beliefs resembles level-2 thinking. Rothschild *et al.* (2012) have taken the approach further and applied level-k thinking to the ARA approach. Their methodology has drawbacks, though, because the level-k approach requires additional assumptions and the problems become quickly intractable due to their growing complexity. The greatest advantage of level-k thinking is that it shows how the level of adversary's thinking affects the optimal decision.

Caswell *et al.* (2011) present a model in which the decision process is evaluated using a Bayesian network with an embedded semi-Markov decision process. Compared to the ARA approach, their model can be used to present the adversary's decision process with greater accuracy. However, as in any decision analysis model, the results are only as good as inputs, and a detailed description of the adversary's thought process would require detailed information about the adversary's resources, values and goals.

Zuckerman *et al.* (2012) represent adversarial activity with a Beliefs-Desires-Intentions (BDI) based model; such models are commonly used to describe teamwork and cooperation. In this approach, the adversary can be modelled as a more nuanced rational agent instead of an omniscient utility maximizer. Yet, the model is not very elaborate, and it can be applied only in zero-sum games in which the goals are easily decomposable.

Applying ARA to military combat modeling

A significant proportion of ARA literature is focused on preventing terrorist threats and, more specifically, on how limited resources should be allocated to combat such threats (*cf.* Pat-Cornell & Guikema 2002, Kardes & Hall 2005, Zhuang & Bier 2007, Golany *et al.* 2009, and Kroshl *et al.* 2015). Nevertheless, in this section we focus more on how ARA can be applied to to military combat modeling and modeling processes, because resource allocation is well covered in earlier research.

We have chosen to examine what possibilities ARA offers for simulating longer chains of events and military deceit, because these are some of the more difficult problems to be handled with existing simulation and analysis tools. To some extent, these topics are interconnected, because deception and misinformation can have major impacts on what happens in the battlefield. Many of the following ideas are still untested, and they are presented as suggestions for worthy topics for future research.

Simulating larger chains of events

The ARA methodology can also be applied to model military operations that are too large to simulate as a single scenario. The scale can become an issue if the number of units involved is too large, or the operation takes place over such a long timeframe that the number of possible paths based on the events becomes excessive. Kangas and Lappi (2006) present how methods of probabilistic risk analysis can be used in conjunction with stochastic combat modeling to analyze longer chains of events. The ARA approach can be used to build on such results to take the analysis one step further. In addition to predicting the success chances of larger operations, it would be possible to predict which ones out of adversary's alternative actions can affect the chain of events most.

Furthermore, ARA can be used to expand a small scale simulation model to a larger scale optimization model. In practice this could mean, for example, using a platform level simulator that can model an aerial battle between fighter aircraft in conjunction with ARA to forecast which decisions would most likely lead to air superiority in the conflict. This approach is not even restricted to using a single simulation model. It would not be significantly more difficult to combine the simulation results from several different models.

Practically any combat model can be used with ARA methodology on condition that the probabilities for each side winning the battle as well as the expected losses on both sides can be calculated. This includes essentially all stochastic combat models and even some deterministic ones. The selection of the combat model must fit the problem at hand. Sometimes the best choice is a platform level Monte Carlo simulation, and sometimes it can be a high level attrition model like the FATHM (Fast Theater Model) (Brown & Washburn, 2000).

In some cases, it is possible to use ARA to model longer chains of events without having to rely on an actual stochastic combat modeling software like Sandis (cf. Kangas and Lappi, 2006). There are also alternative, lighter stochastic computational models that can be used to predict the outcome of a duel between two platoon sized forces (Lappi *et al.,* 2012; Åkesson, 2012; Roponen, 2013). These models can be used to significantly cut down the time for calculating all the success probabilities and the expected losses in different stages of the chain. There are also additional time savings from not having to create a complete model scenario, which, as noted earlier, is a time consuming process. The use of the lighter duel simulation methods could be automatized to a certain degree, because they require far fewer input parameters.

To offer a rough outline for how ARA can be applied to a longer chain of events, we present an example case in which modeling the events as a single combat simulation would likely be very time consuming and would not offer any real benefits over the ARA approach. That is, we examine a situation in which there are two bridges that can be used to cross a river. One party wants to cross the river and the other wants to prevent that from happening by using military force. The Defender has three platoons of soldiers available for defending the bridges. The Attacker has six platoons which try to get across.

In order to examine a problem such as this, one first needs to a narrative of what chains of events are possible as a result of the decisions that the actors can take. To keep the problem tractable, the number of possible decision options as well as the end results of random events needs to be kept to a minimum, because the number of calculations required grows exponentially with each step. The problem has been presented as an influence diagram in Figure 4.

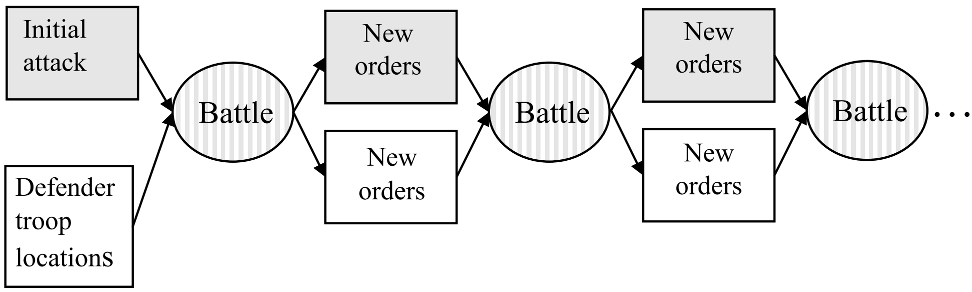

Figure 4: Influence diagram depicting a longer chain of events. The darker nodes are associated with the Attacker, as perceived by the Defender, as uncertainty nodes and vice versa. The striped nodes are uncertain events common to both sides.

The Defender's first decision is how he should deploy his troops initially. This includes deciding how he should divide his troops between the two bridges and reserves, but depending on the level of sophistication of the combat model used for the calculations other variables like the use of terrain, fortifications and mines could be included as well. The Attacker also needs to decide how he is going to use his troops for the initial attack. This includes the number of troops used and the target (or many targets) of the attack. Depending on the modeled situation the Attacker might be operating with very limited information. The initial attack might be used to just scout the Defender's strength.

The influence diagram in Figure 4 shows the two decisions as independent, but it does not really need to be so. Depending on which situation we wish to model, either the Attacker or the Defender could act after finding out what the other is planning to do. It is even possible to account for asymmetric information on both sides, arising for instance from greater familiarity with the terrain, military deception or some other reason. For the sake modelling parsimony, the number of possible choices should be kept as low as possible. This means, for example, that the troops should be deployed as whole platoons instead of single soldiers.

Once the possible initial decisions have been determined, the probabilities of all the possible results of the first battle for all the possible initial conditions need to be computed. The number of possible battle results can become computationally overwhelming, because all possible combinations of casualties on both sides are technically different outcomes. It is not feasible to calculate the next step if the number of possibilities is in the hundreds. The amount can be reduced to a more manageable number by assuming that the losses suffered by each side are conditional only on the result of the battle and not on each other. In this way, the possible results can be limited to wins and losses or some other smaller set. One way to choose the set of possible results is to take a look at the next step and identify which results would lead to different decisions and use those. Depending on the modeled chain of events, this can be easy, nearly impossible, or anything in between these extremes.

When the possible results of the first battle have been determined, we proceed to the decisions that follow it, then to the next battle, and so forth. In some these outcomes, either the Attacker or the Defender will not be able to continue effectively, which shortens the chain of events. For the sake of tractability, one also has to decide where these chains of events will have terminated so that there is no more fighting. It is possible that in more complex cases, the influence diagram is not ideal for visualizing the problem. Alternatively, the problem can be shown as a decision tree, because the ARA methodology is not tied to influence diagrams.

After all the possible chains of events have been elaborated and the corresponding possible end results have been determined, we proceed by calculating backwards from the end to estimate the probabilities of these results. Towards this end, we first solve the ARA problem formed by the last decisions in the chain and the ensuing battle. To find out the probability that the initial conditions for those ARA problems are met, we then solve the ARA problems formed by the battle and decisions preceding them and so forth until the first decisions have been analyzed. The utility function for each problem in the chain is formed from the maximized expected utility gained from each outcome as given by solving the ARA problems following it.

Modeling the effectiveness of military deceit

"All warfare is based on deception." (Sun Tzu) Using deceit to gain an upper hand against an adversary is an absolutely integral part of military tactics and strategy. Still, the effects of deceit are very difficult to predict and simulate with existing operational analysis and combat modeling software. Because the effects cannot be readily reduced to mathematical formulas, modeling the effects of deceit relies usually on expert judgements. In the context of combat modeling, this usually means that the required expert opinions are provided by the operator of the simulation tool.

A common alternative is to use wargames to model the uncertainties in human decision making; but even this approach also has problems as wargames ignore many aspects of reality. Questions of solvability do not arise in wargames, because the aim is not to determine optimal tactics. Rather, realizing that wars are fought by humans, wargames study the decision process of humans. One problem in this approach is that in a game, the player can make decisions that he would not really make as long as these decisions produce good results in the simulation. For example, in a simulated environment casualties may not have the weight that they would have in actual combat. The second problem is that wargames often capture typical decision making behavior (rather than optimal decisions) because the players play a small number of games only. Thus, for instance, the resulting lack of repetition may overstate the effectiveness of new weapon systems, because the opponent does not have time to learn and adapt his tactics to counter these systems. To some extent, the lack of repetition may be deliberate due to the fear that the players would learn to use the artificialities of the wargame to their advantage instead of developing better military strategies. Another reason for the lack of repetitions is that wargaming is time consuming and expensive. (Washburn & Kress, 2009, 111-130)

The ARA approach could be used to assess the effects of deceit tactics on the decision making of the adversary. Specifically, the ability to model the effects of the adversary's altered perceptions would be a useful complement to the elicitation of expert judgements. Mathematical equations are, after all, immune to effects of optimistic thinking.

ARA can be used relatively easily to model situations in which the adversary is deliberately misinformed about the strength or capabilities of the opponent, for instance as a result of hiding troop movements and employing dummy units or decoys. Then, ARA helps estimate the effect the misinformation on the adversary's decision making and whether this effect is beneficial so that the benefits outweigh the costs. In the next section, we give an example of such an estimation process.

The ARA approach is not limited to deceit that happens on the battlefield; indeed, military deceit is pervasive in military planning and decision making. In theory, ARA can be used to study the effects of almost any misinformation, but usually these effects can not be modeled with combat simulation models (except in the case of misinformation that relates to the number or capabilities of weapon systems, sensor systems or military units). While the ARA methodology does not give tools for predicting the probability with which deception will succeeds, it helps assess the possible effects of successful deception may be, which helps decide what information is worth hiding or altering.

Example of applying ARA to a military deceit problem

To demonstrate how ARA can be applied, we use it here to analyze a relatively simple tactical problem which illustrates some of the key aspects of the approach.

The problem

Consider the following situation in which there are two adversaries: the Defender and the Attacker. The Defender has two valuable targets he needs to protect, Target 1 and Target 2. The Defender has a total of 60 troops, 40 of whom are situated at Target 1 and 20 at Target 2. The Defender can try to secretly move troops from one target to another, but there is a possibility that the Attacker will notice the troop movement. After the Defender has moved the troops he wants to move, the Attacker will decide which target he will attack. If the Defender succeeds in moving the troops without the Attacker noticing, the Attacker will have to decide his target using incomplete information. We solve this problem from the point of view of the Defender. Figure 5 shows the influence diagram from the Defender's point of view.

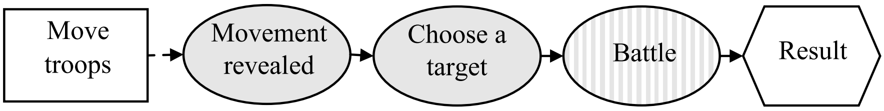

Figure 5: Influence diagram of the example case from the point of view of the Defender.

We denote the set of options the Defender as D, which represents number of squads of five soldiers moved from Target one to Target 2, and the options of the Attacker $A = \{1,2\}$ which represents the choice between the two targets. The uncertainties in this case are whether the troop movement is revealed to the attacker R ($R = 1$ if it is and $R = 0$ if not) and the outcome of the battle S ($S = 1$ if the Defender wins and $S = 0$ if the Defender loses). The utility functions over the costs are $u_D(c_D, c_A)$ and $u_A(c_D, c_A)$, with costs dependent on the actions of the Defender and the Attacker.

The Defender wants to solve

$$\psi = \max_{d \in D} \left[\psi(d) = \sum_{a \in A} \sum_{r \in R} [p_D(r) p_D(a|d,r) \psi_D(a,d)] \right]. \quad (3)$$

In order to solve this problem, it is necessary to assess the probabilities over the costs, conditional on (a, d, S); and about S, conditional on (a, d). In this case, the Attacker and the Defender have different assessments: for example, for success, these are $p_D(S = 1|a, d)$ and $p_A(S = 1|a, d)$, respectively. It is likely that the Attacker's assessment of the success of the assault differs from that of the Defender, because the Attacker may not know the Defender's decisions (see Figure 6). The expected utility for the Attacker, resulting from (a, d) is

$$\begin{aligned} \psi_A(a,d) = p_A(S = 0|a,d) \sum_{c_A} \sum_{c_D} [u_A(c_A, c_D, a, S = 0) \pi_A(c_A, c_D|a, d, S = 0)] \\ + p_A(S = 1|a,d) \sum_{c_A} \sum_{c_D} [u_A(c_A, c_D, a, S = 1) \pi_A(c_A, c_D|a, d, S = 1)]. \end{aligned} \quad (4)$$

The Defender's expected utility is similar.

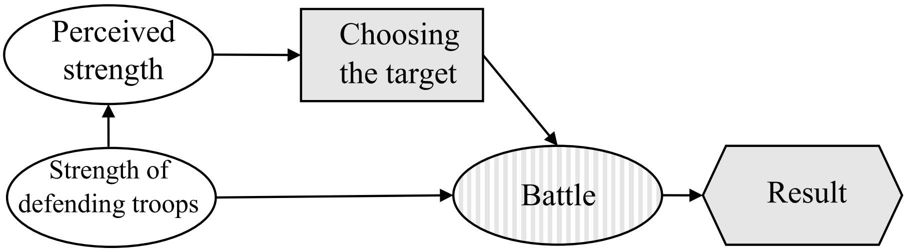

Figure 6: Influence diagram of the example case from the point of view of the Attacker.

We now solve the game from the Defender's point of view. The Defender has 40 men defending Target 1 and 20 men defending Target 2. The Defender considers value of Target 1 to be 50 and value of Target 2 to be 30. He considers the value of a single combat ready soldier to be 1. He has the option moving troops between the targets without the Attacker's knowledge. There is, however, a probability p_R that the attacker finds out about the troop movement. The Defender also estimates that the Attacker has at least 20 men but no more than 35, and he thinks that the most likely number is 30, so he fits a triangular distribution.

Using the strength estimates of both forces, the Defender can use, for example, a stochastic combat model to calculate p_D and π_D. Specifically, the Defender assesses that the utility in this situation follows the function

$$u_D(S, a, c_A, c_D) = \begin{cases} 0.1c_A - c_D, & if\ S = 1 \\ -50 + 0.1c_A - c_{D,} & if\ S = 0\ and\ a = 1 \\ -30 + 0.1c_A - c_D & if\ S = 0\ and\ a = 2, \end{cases} \quad (5)$$

where $S = 1$ corresponds to the situation where the Defender manages to hold the target area, $a = 1$ corresponds to the situation where the Attacker decides to attack Target 1 and $a = 2$ corresponds to the Attacker deciding on Target 2, and c_A and c_D are the Attacker's and the Defender's losses respectively.

Solving the problem for the Defender is not sufficient for determining p_D, π_D and u_D. To calculate the expected utility from decision *d,* he first needs to estimate $p_D(a|d, r)$. Towards this end, the Defender needs to solve the problem from the viewpoint of the Attacker. He assumes that the Attacker, too, is an expected utility maximizer. The problem is presented from the Attacker's point of view in Figure 6.

The Defender estimates that the Attacker thinks the Defender has 36 to 44 men at Target 1 and 18 to 22 men at Target 2 (with all values equally probable), and the Attacker has probability $p_D(R = 1) = 0.1$ of finding out about the Defender's troop movement. If the Attacker detects the Defender moving troops, he will be able to accurately count the number of troops moved. Using those strengths for his estimates he can use the same stochastic combat model used to solve p_D and π_D to calculate p_A and π_A.

The Defender estimates that the Attacker's utility function is similar to his own. However, the Defender does not know for sure how valuable each target is to the Attacker. He models this uncertainty by adjusting the weights of successful capture of each target S. Thus, he estimates that the Attacker's utility function is

$$u_A(S, a, c_A, c_D) = \begin{cases} 0.1c_D - c_D, & if\ S = 1 \\ 50 + U_1 + 0.1c_D - c_{A,} & if\ S = 0\ and\ a = 1 \\ 30 + U_2 + 0.1c_D - c_A & if\ S = 0\ and\ a = 2, \end{cases} \tag{6}$$

where the distributions U_1 and U_2 are uniformly distributed over the interval $[-5, 5]$.

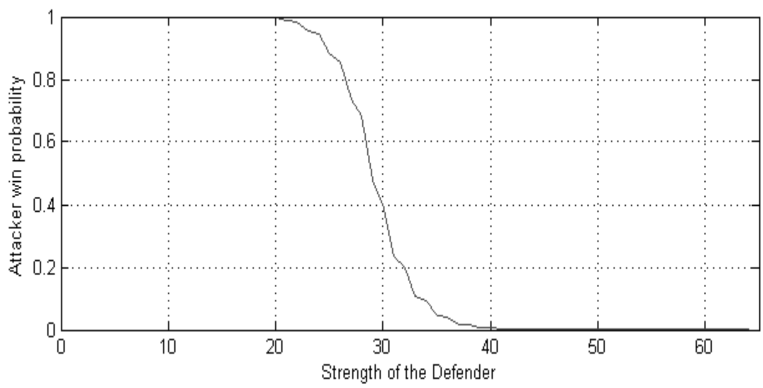

Figure 7: The Attacker's win probability with strength of 30 as a function of the Defender's strength. The minor perturbations in the curve are caused by the assumption that the battle is lost when the unit has lost half of its troops, which means that even strengths are slightly less advantageous for the Defender.

We solve the problem step by step. First, the Defender will:

1. Calculate the success probabilities and expected losses for both sides for all the possible combinations of strengths of both sides (Figure 7).
2. Calculate the Attacker's expected utilities ψ_A for attacking each target for all possible strengths of the Attacker's force taking into account the uncertainties with u_A the fact that the Attacker has knowledge of the Defender's troop movement on probability $p_D(R = 1)$.
3. Compare the expected utilities to get an estimate for the probability of an attack on each target for each possible strength of the Attacker.
4. Consider the probability of an attack with a specific strength of the attacker and the probability for each of those strengths to calculate $p_D(a|d, r)$.
5. Calculate $\psi_D(a, d)$.
6. Use $p_D(a|d, R)$ and $\psi_D(a, d)$ to determine the decision *d* which maximizes his expected utility.

We used the approximative method in Roponen (2013) to simulate a duel between two forces in order to calculate the probabilities in (p_D, π_D, p_A and π_A) in step 1, because the details for this method are publicly available and it produces the results of battles between two infantry units very efficiently. This program code was used to examine the win probabilities for all possible strengths of both sides. Situations in which a tie was predicted were counted as the Defender's victory, because the Attacker would be unable to capture the target. Moreover, a unit was assumed to lose the battle if it lost half of its fighting strength, which is the cause behind the roughness of curves depicted in Figures 7 and 8, because the unit strengths are non-negative integers and odd numbers are not divisible by two.

We then wrote a program code to search the remaining steps exhaustively, to calculate the expected utilities, and to determine the attack probabilities $p_D(a|d)$. Because the Attacker perceived that target 2 was significantly weaker, he always chose to attack this target unless he found out about the troop movement, in which case he chose the target with actually higher expected utility. The attack probability on target 1 thus varied between 0 and 0.1. Then the expected utilities of the Defender were calculated from p_D and π_D as seen in Figure 8.

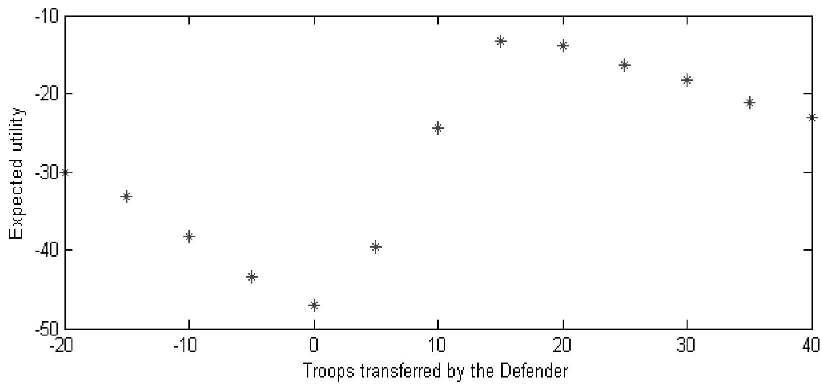

Figure 8: The expected utility gained by the Defender as a function of troops transferred from target 1 to target 2.

The highest expected utility gained by the Defender, $\psi_D \approx -13.4$, was achieved by transferring three squads or 15 soldiers from target 1 to target 2. This gave the Defender 0.933 probability of winning the battle if the Attacker chose to attack target 2 (which would happen if he does not notice the troop movement) and 0.281 probability of winning if the Attacker chose to attack target 1.

Future research

Adversarial risk analysis (ARA) is a relatively new research area which is becoming more prominent in the context of counter-terrorism and corporate competition. In this paper, we have discussed and illustrated the application of ARA to the modelling of longer chains of events and the effects of military deceit. There are also many other possible ways of applying the ARA methodology to combat modeling.

Arguably the most important reason for military combat modeling is that it provides support for strategic, tactical or technical decision making (Tolk, 2012, 55-78). However, it is not straightforward to translate the results of combat models into decision recommendations (Davis & Blumenthal, 1991). One accessible way of using ARA in the context of simulation is to perform an exhaustive portfolio analysis of all relevant strategies. Such methods have been used to assess the cost-efficiency of different combinations of weapons systems, for example (Kangaspunta *et al.*, 2012).

In the same vein, ARA could be used to predict the most likely responses of the adversary and to calculate the expected utilities of each strategy under different conditions. The applicability of this approach depends how well the process can be streamlined and automated so that informative results about strategies can be provided more quickly than through manual analyses of combat simulations.

To enrich the possibilities of using combat simulation models, ARA could also be used to reduce the need for user interaction within existing combat modeling tools. Simple adversarial intent models have been used in professional wargaming to simulate intelligent forces (Santos & Zhao 2006). This notwithstanding, most combat models do not yet include algorithms that would represent the human thought processes involved in tactical or strategic decisions (Washburn & Kress, 2009, 111-130). Depending on the model, practically all higher level decisions concerning the position and strength of forces are made by the operator. As a result, the time required to create a scenario is usually significantly longer than the time required to calculate the results (Lappi, 2012).

The development of simulation models in which the units are able to make simple tactically sensible decisions would widen the range of problems that can be analyzed by using approaches such as data farming. The ARA methodology could then be used as a basis for these kinds of algorithm. Here, ARA has advantages over using game theory, because it accounts for uncertainties and even misinformation.

However, ARA cannot be readily applied to very low level or continuous decision making, because the required calculations would become just too overwhelming. ARA can be used most effectively in situations where the attention can be restricted to choices among relatively few possible strategies. If need be, it may be possible to simplify the problem by restricting attention to some plausible chains of events instead of calculating all possible chains of events. Analogues approaches towards simplification have been employed in constructing artificial intelligence systems for games such as Go and Chess, and they have been applied even in video game AI development (Churchill *et al.* 2012).

Conclusions

The ARA methodology has already found many uses in analyzing counter terrorism and corporate competition (Rios Insua *et al.,* 2009). In this paper, we have discussed the relevance of this methodology to military combat modeling and presented concrete examples of how it can be applied. Specifically, we have outlined possible uses for the ARA approach in the context of modeling deceit and using ARA to stretch the limits of existing combat models to model longer chains of events. Many of these ideas are tentative and call for more research before they can be implemented into existing simulation. Another challenge is that real battles are extremely complex, they involve thousands of decisions, and the goals and resources are highly uncertain. Still, by focusing on the most important decision situations and decision alternatives can provide valuable insights.

We also presented an illustrative example in which ARA was combined with stochastic combat modeling to calculate the effects of military deceit. In this example, most calculations for solving the ARA part of the model were relatively straightforward and could be implemented into software code (there are numerous tools for calculating the results of battles; see, Kangas, 2005). We therefore believe that it possible to develop software tools for considerably more complex problems in which the dependencies between the adversaries' utilities and their decision behaviour are explicitly modelled. More generally, there is much potential in using the ARA approach to tackle realistic problems through stochastic combat modelling. This would serve to push the boundaries of ARA modelling in an important application area.

Fundamentally, ARA has much to offer for military combat modeling, because it is able to combine the conventional statistical approach of risk analysis–which is already widely employed in combat modeling–with fresh game-theoretical perspectives that help predict what one's opponents are likely to do. ARA can also be used to build optimization models on top of existing simulation models, which gives possibilities for new uses for these simulation models as well as new ways of visualizing the results of such models.

References

Banks, D., Petralia, F., & Wang, S. (2011). Adversarial risk analysis: Borel games. *Applied Stochastic Models in Business and Industry* 27(2), pp. 72-86.

Brown, G. & Washburn, A.R. (2000 rev. 2004). *The fast theater model (FATHM),* Project Report, (NPS-OR-01-002-PR), Naval Postgraduate School, Monterey, CA.

Camerer, C. (2003). *Behavioral Game Ttheory: Experiments in Strategic Interaction*. Princeton University Press, New Jersey.

Caswell, D. J., Howard, R. A., & Paté-Cornell, M. E. (2011). Analysis of national strategies to counter a country's nuclear weapons program. *Decision Analysis*, *8*(1), 30-45.

Churchill, D., Saffidine, A., & Buro, M. (2012). Fast Heuristic Search for RTS Game Combat Scenarios. In *AIIDE.*

Davis, P. K., & Blumenthal, D. (1991). *The base of sand problem: A white paper on the state of military combat modeling* (No. RAND/N-3148-OSD/DARPA). Defense Advanced Research Projects Agency, Arlington VA.

Golany, B., Kaplan, E. H., Marmur, A., & Rothblum, U. G. (2009). Nature plays with dice–terrorists do not: Allocating resources to counter strategic versus probabilistic risks. *European Journal of Operational Research*, *192*(1), 198-208.

Howard, R. A. & Matheson, J. E. (2005). Influence diagrams. *Decision Analysis*, *2*(3), pp. 127-143.

Kadane, J. B. & Larkey, P. D. (1982). Subjective probability and the theory of games. *Management Science*, *28*(2), pp. 113-120.

Kangas, L. (2005). Taistelun stokastinen mallinnus. Master's Thesis, Helsinki University of Technology. http://sal.aalto.fi/publications/pdf-files/tkan05.pdf Accessed: 2014-05-07

Kangas, L. & Lappi, E. (2006) Probabilistic risk analysis in combat modeling. In:

Hämäläinen, J. (ed.) *Lanchester and Beyond. A Workshop on Operational Analysis Methodology.* PVTT Publications 11

Kangaspunta, J., Liesiö, J., & Salo, A. (2012). Cost-efficiency analysis of weapon system portfolios. *European Journal of Operational Research*, 223(1), pp. 264-275.

Kardes, E., & Hall, R. (2005). Survey of literature on strategic decision making in the presence of adversaries. *Unpublished report.*

Kovenock, D., Roberson, B. (2010). Conflicts with multiple battlefields, *CESifo working paper: Empirical and Theoretical Methods, No. 3165*

Kroshl, W. M., Sarkani, S., & Mazzuchi, T. A. (2015). Efficient allocation of resources for defense of spatially distributed networks using agent-based simulation. *Risk Analysis* (to appear).

Lappi, E. (2012). Computational methods for tactical simulations. *Julkaisusarja 1. N:o 1/2012*. Doctoral Thesis, National Defence University, Finland

Lappi, E., Pakkanen, M., & Åkesson, B. (2012). An approximative method of simulating a duel. In: *Proceedings of the Winter Simulation Conference*, *WSC '12*, pp. 208:1-208:10

Lappi, E., Pentti, J., Åkesson, B., Roponen, J., Valtonen, J., Koskinen, J., Burhan, U., Sivertun, Å., and Hämäläinen, J. (2015). Team 4: Data farm. *manuscript.*

Meng, S., Wiens, M., & Schultmann, F. (2014). A game-theoretic approach to assess adversarial risks. In: Brebbia, C.A. (ed.), *Risk Analysis IX*, WIT Press, Southampton, UK, pp. 141-152.

Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA.

Paté-Cornell, E., & Guikema, S. (2002). Probabilistic modeling of terrorist threats: A systems analysis approach to setting priorities among countermeasures. *Military Operations Research*, 7(4), pp. 5-23.

Reese, W. (1980). Deception in a game theoretic framework. In: Daniel, D. C., Herbig, K. L., Reese, W., Heuer, R. J., & Sarbin, T. R. (1980). *Multidisciplinary Perspectives on Military Deception* (No. NPS-56-80-012A). Naval Postgraduate School, Monterey, CA.

Rios Insua, D., Rios, J., & Banks, D. (2009). Adversarial risk analysis. *Journal of the American Statistical Association*, 104(486), pp. 841-854.

Rios, J., & Rios Insua, D. (2012). Adversarial risk analysis for counterterrorism modeling. *Risk analysis*, *32*(5), pp. 894-915.

Roponen, J. (2013). Kaksintaistelun approksimatiivinen mallintaminen. Bachelor’s Thesis, Aalto University. http://sal.aalto.fi/publications/pdf-files/trop13_public.pdf Accessed: 2014-05-07

Rothschild, C., McLay, L., & Guikema, S. (2012). Adversarial risk analysis with incomplete information: A level-k approach. *Risk Analysis*, 32(7), pp. 1219-1231

Santos Jr, E., & Zhao, Q. (2006). Adversarial models for opponent intent inferencing. In: Kott A., & McEneaney W. (2006) *Adversarial Reasoning: Computational Approaches to Reading the Opponents Mind.* Chapman & Hall/CRC, Boca Raton, 1-22.

Tolk, A. (2012). *Engineering Principles of Combat Modeling and Distributed Simulation*. Wiley. USA.

Tzu, S. (2004 print) *The art of War.* Trans. Giles L. Project Gutenberg.

Washburn, A. R., & Kress, M. (2009). *Combat Modeling*. Springer. Heidelberg.

Zhuang, J., & Bier, V. M. (2007). Balancing terrorism and natural disasters-defensive strategy with endogenous attacker effort. *Operations Research*, 55(5), pp. 976-991.

Zuckerman, I., Kraus, S., & Rosenschein, J. S. (2012). The adversarial activity model for bounded rational agents. *Autonomous Agents and Multi-Agent Systems*, *24*(3), 374-409.

Åkesson, A. (2012). Automatic calculation of win probabilities and conditional strength distribution of units in stochastic simulation model. Master's Thesis, Åbo Akademi University.

Publication II

Juho Roponen, David Ríos Insua, and Ahti Salo. Adversarial risk analysis under partial information. *European Journal of Operational Research*, 287(1), 306-316, November 2020.

European Journal of Operational Research 287 (2020) 306–316

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

Adversarial risk analysis under partial information☆

Juho Roponen[a,*], David Ríos Insua[b], Ahti Salo[a]

[a] Systems Analysis Laboratory, Department of Mathematics and Systems Analysis Aalto University School of Science, PO Box 11100 Aalto FI 00076, Finland
[b] ICMAT–Institute of Mathematical Sciences, Nicols Cabrera, n 13–15 28049 Madrid, Spain

A R T I C L E I N F O

Article history:
Received 18 February 2019
Accepted 20 April 2020
Available online 28 April 2020

Keywords:
Risk analysis
Decision analysis
Game theory
Stochastic dominance
Combat modeling

A B S T R A C T

Adversarial risk analysis provides one-sided decision support to decision makers faced with risks due to the actions of other parties who act in their own interest. It is therefore relevant for the management of security risks, because the likely actions of the adversary can, to some extent, be forecast by formulating and solving decision models which explicitly capture the adversary's objectives, actions, and beliefs. Yet, while the development of these decision models sets adversarial risk analysis apart from other approaches, the exact specification of the adversary's decision model can pose challenges. In response to this recognition, and with the aim of facilitating the use of adversarial risk analysis when the parameters of the decision model are not completely known, we develop methods for characterizing the adversary's likely actions based on concepts of partial information, stochastic dominance and decision rules. Furthermore, we consider situations in which information about the beliefs and preferences of all parties may be incomplete. We illustrate our contributions with a realistic case study of military planning in which the Defender seeks to protect a supply company from the Attacker who uses unmanned aerial vehicles for surveillance and the acquisition of artillery targets.

1. Introduction

As a growing field of research, Adversarial risk analysis (ARA) (Banks, Aliaga, & Ríos Insua, 2015) provides one-sided decision support to a decision maker who is faced with risks that depend on the decisions of other self-interested parties. This makes ARA relevant for the analysis and management of security risks, including those of terrorism, military operations and cyber threats. There are numerous reported ARA applications in areas such as anti-piracy (Sevillano, Ríos Insua, & Rios, 2012), counter terrorism (Ríos & Insua, 2012), combat modeling (Roponen & Salo, 2015), and anti-IED war (Wang & Banks, 2011), for instance.

Specifically, ARA helps characterize the likely actions of all parties by building and analyzing multi-agent representations of the decision problem, taking into account their values, objectives, goals, capabilities and beliefs of the parties involved. Particular attention is paid to modeling the information on the basis of which the parties, referred to as *adversaries*, make their decisions. The aim is to build realistic models which, unlike most game theoretical analyses (Antos & Pfeffer, 2010; Ozdaglar & Menache, 2011), do not necessitate far-reaching and partly unrealistic assumptions about common knowledge which shared by all parties.

While the modeling of the adversaries' decision processes makes ARA a powerful approach, there are notable challenges, too. In particular, it can be difficult to produce accurate estimates about how the adversary's preferences and beliefs evolve over time. This could be the case, for instance, in situations where two adversaries, the Defender ('she') and the Attacker ('he'), of which the Defender first chooses what countermeasures she will adopt in her defence, whereafter the Attacker, knowing the Defender's choice, updates his beliefs and proceeds by deciding how to attack (Xu & Zhuang, 2016). These difficulties notwithstanding, attempts to building realistic representations of the intertwined decision problem should be made, because such representations can yield valuable insights and because the reliance on overly simplistic models of the adversary's preferences will lead to sub-optimal countermeasures (Nikoofal & Zhuang, 2015).

☆ This research has been financially supported by MATINE - The Scientific Advisory Board for Defence of the Finnish Ministry of Defence. David Rios Insua acknowledges the Spanish Ministry of Economy and Innovation program MTM2014-56949-C3-1-R, MTM 2017-86875-C3-1-R, the AXA-ICMAT Chair on Adversarial Risk Analysis, the European Union's H2020 Program for Research, Technological Development and Demonstration, under grant agreement no. 740920 (CYBECO), and the Visiting Researcher programme of the Aalto Science Institute.

* Corresponding author.
E-mail addresses: juho.roponen@aalto.fi (J. Roponen), david.rios@icmat.es (D. Ríos Insua), ahti.salo@aalto.fi (A. Salo).

https://doi.org/10.1016/j.ejor.2020.04.037
0377-2217/

In this paper, we develop conceptual, mathematical and computational methods based on the concepts of partial information and stochastic dominance to identify which of the adversaries' decisions are non-dominated and therefore plausible (Levy, 1992). The practical relevant of these methods stems from the fact that the set of these non-dominated decisions tends to be much smaller than the set of all possible decisions. Thus, they provide useful decision support even in situations where complete information about the adversaries' utilities and/or beliefs cannot be obtained. Specifically, we allow for the possibility that there is only partial information about the preferences and beliefs of the Attacker (or even about the the Defender), operationalized through subsets of probability distributions and multi-attribute utility functions. As it turns out, this approach is quite flexible and suitable for providing informative decision support. We note that what we call partial information differs from what is understood by Rothschild, McLay, and Guikema (2012) who examine situations where the Attacker can only partially observe the Defender's defence decision.

We illustrate our methods in the light of a sequential defend-attack game, as such games are particularly important in the security domain (Brown, Carlyle, Salmerón, & Wood, 2006; Zhuang & Bier, 2007). Previously, stochastic dominance has been applied to examine both simultaneous or continuous games, see Fishburn (1978) and Rass, König, and Schauer (2017). However, as noted by Fishburn (1978), simultaneous games do not always have realistic properties and consequently simple game theoretic solutions may not provide meaningful decision support. For example, in the continuous game considered by Rass et al. (2017), the Attacker's decision problem is modeledso that the Attacker only seeks to cause maximal harm to the Defender, which effectively reduces the problem to a two-player zero sum game.

For completeness, we first review the Bayesian Nash equilibrium solution which is used in (i) methods which rely on common prior probability distributions over players' types and (ii) the standard ARA solution which uses probability distributions to model uncertainties about the adversary's preferences and beliefs. Then, we present our key methodological contributions in two variants, starting from the situation in which there is incomplete information only about the preferences and beliefs of the adversary (in our case the Attacker), modeled though sets of utility functions and probabilities. We then consider the situation in which the information about the Defender's (own) preferences and beliefs, too, may be incomplete, using stochastic dominance to produce meaningful analyses (Shaked & Shanthikumar, 2007). Finally, we show how the partial information approach can be extended to analyze sequential games in so-called regular influence diagrams.

We also present a realistic case study to illustrate how our methods can be applied to different kinds of sequential problems. In our case study, the Defender seeks to determine an efficient portfolio of countermeasures to protect a supply company against the Attacker's unmanned aerial vehicle (UAV) reconnaissance and the subsequent specification of artillery targets. The Attacker, in turn, seeks to choose UAV and artillery systems which are cost-efficient in responding to the Defender's countermeasures. We also discuss other security and military problems which are amenable to our methods.

2. Bayesian models

Adversarial problems are often modeled as two-player games (Cox, Jr, & Anthony, 2009; Washburn, Kress et al., 2009). Real world multi-agent decision problems can often be modeled as games of *incomplete information*, meaning that there are some players who do not know all the rules of the game, such as the capabilities, utilities and decision processes of the other players. In *stochastic* games there are uncertain chance events, such as weather or the outcome of a military combat. Furthermore, in games of *imperfect information*, there are players who cannot observe the previous actions of the other players and/or the outcomes of random events before it is their turn to act.

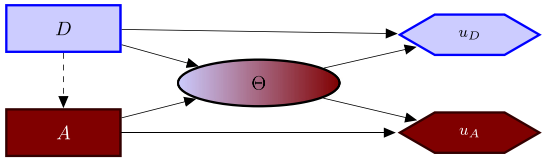

Fig. 1. A bi-agent influence diagram depicting the sequential defend-attack game.

By far the most common approach to solving games of incomplete information is based on the Bayesian Nash equilibrium. Modeling the game as what is commonly known as a Bayesian game transforms the game of incomplete information into a stochastic game of imperfect information. In a Bayesian game, the type of each player is defined by his/her utilities, beliefs, decision alternatives, available resources, etc. Every player can observe his/her own type but not the types of other players. (Harsanyi, 1967)

In order to solve the stochastic game, it is necessary the make assumptions about the players' types as well as their beliefs concerning each others' types, formalized through prior probability distributions. In the traditional game theoretical approach, the players are assumed to have common knowledge about these prior distributions (Bier, Oliveros, & Samuelson, 2007). It is even possible to derive the prior distributions based on a level-k approach (Rothschild et al., 2012). While choosing the right probability distributions is important and has both practical and philosophical implications, the question of where these distributions come from does not affect how the game is solved.

In ARA applications, in contrast, the prior distributions are based on the subjective estimates of the player whose decisions are being supported and who seeks to choose the best decision alternative(s) under adversarial uncertainty (Ríos Insua, Rios, & Banks, 2009).

2.1. Influence diagrams

We visualize ARA problems using bi-agent (or even multi-agent) influence diagrams (Banks et al., 2015). In an influence diagram, rectangular nodes indicate decisions, circular nodes depict uncertain chance events, and hexagonal nodes correspond to utilities. Directed arcs are employed to connect the nodes so that continuous arcs represent probabilistic dependencies and dashed arcs represent the information that is available at decision nodes. Based on information about (i) the adversaries, (ii) the order in which they make decisions; (iii) the available decision alternatives at each decision node; (iv) the information available to the adversary at each decision node; and (v) the utilities for the adversaries resulting from any sequence of events, it is possible to produce plausible predictions concerning the adversaries' likely actions, ensuing outcomes, and consequences for each adversary.

Our methodological development refers extensively to the basic sequential defend-attack game, visualized as a bi-agent influence diagram in Fig. 1. This is a stochastic game of incomplete information in which all the actions and random events are observable after-the-fact, so it also features perfect information. The Defender's nodes are shown in blue and those of the Attacker in red. The node Θ represents an uncertain outcome which is common to both adversaries and colored in both colors. The dashed arc between the Defender's and Attacker's decisions describes an information structure, indicating that the Attacker knows the Defender's decision when the Attacker (A, referred to as 'he') makes his decision.

Specifically, the Defender (D, referred to as 'she') first implements her defensive decision $d \in \mathcal{D}$. Next, the Attacker observes this defense and makes his decision $a \in \mathcal{A}$. Finally, the combination of these decisions affects the probabilities of possible outcomes at the chance node Θ. The realized outcome $\theta = \Theta(a, d)$ is typically uncertain and thus there is uncertainty about the consequences $c(d, a, \theta)$. Often, these consequences involve multiple attributes, such as casualties, material losses and monetary costs for both adversaries.

The Defender evaluates the consequences with her utility function $u_D(c(d, a, \theta))$ and the Attacker with his utility function $u_A(c(d, a, \theta))$. For brevity, these utilities are denoted by $u_D(d, a, \theta)$ and $u_A(d, a, \theta)$, respectively. In the adversarial setting, the goals and, thus also the utility functions, of the players are different and often opposite.

2.2. Bayesian Nash equilibrium solution

As an example, we first determine the Bayesian–Nash equilibrium for the sequential defend-attack game in Fig. 1. From the Defender's perspective, the uncertainties the Defender faces about the Attacker are modeled through a probability distribution over the set of Attacker's possible types T_A, defined as combinations of $\mathcal{U}_A$, is the set of Attacker's possible utility functions, and $\mathcal{P}_A$, the set of his possible probability estimates for the outcomes θ as a results of decisions a and d. Thus, realizations of the Attacker's type are pairs of utility functions and probabilities $T_A = (u_A, p_A)$. The solution for the game is obtained as follows:

1. At the outcome node Θ, compute the Attackers' expected utilities
$$\psi_A^{u_A, p_A}(d, a) = \int_{\theta \in \Theta} u_A(d, a, \theta) p_A(\theta | d, a) d\theta,$$
for every $(u_A, p_A) \in \mathcal{U}_A \times \mathcal{P}_A$ and decisions $a \in \mathcal{A}$ and $d \in \mathcal{D}$.
2. At the Attacker's decision node A, compute the optimal attacks for the observed defense d using the Attacker's beliefs and utilities
$$A^{*u_A, p_A}(d) = \arg\max_{a \in \mathcal{A}} \psi_A^{u_A, p_A}(d, a),$$
and forecast the attack a through
$$p_D(a|d) = \mathbb{P}(A^*(d) = a),$$
while taking into account the probability $\mathbb{P}(T_A = (u_A, p_A))$.
3. At the Defender node D, compute the optimal defense
$$d^* = \arg\max_{d \in \mathcal{D}} \int_{\mathcal{A}} \psi_D(d, a) p_D(a|d)\, da,$$
where the Defender's expected utility is
$$\psi_D(d, a) = \int_{\Theta} u_D(d, a, \theta) p_D(\theta | d, a) d\theta.$$

Solving the Attacker's decision problem is more straightforward, because the Attacker can observe the defense d before choosing his mode of attack a.

3. Games of partial information

Bayesian Nash equilibrium is the most widely used solution concept for games of incomplete information. Its application assumes that every player solves the game using Bayesian approach, i.e., all players assign subjective probability distributions to the parameters they do not know (Harsanyi, 1967). In practice, this assumption can be problematic, because defining subjective probability distributions over the other player's probability estimates and utilities, or equivalently their type, can pose challenges. It also presumes that all players are willing and able to specify these probability distributions.

In order to support adversarial risk analysis when well-defined probability distributions over players' types cannot be elicited, we explore how a game between rational players can be analyzed *without* such distributions. By rational, we mean that the players seek to maximize their own expected utilities. We also assume that the players have some knowledge about their adversary. Specifically, they are assumed to know that the other player's type, characterized by a combination of utilities and beliefs (u, p), belongs to a subset of possible types $\mathcal{T} \subseteq \mathcal{U} \times \mathcal{P}$. In what follows, this characterization based on set inclusion is referred to as *partial information*.

Using these assumptions, we revisit the Defend-Attack game in Fig. 1. The Defender knows that the Attacker's type $(u_A, p_A) \in \mathcal{T}_A$. For a given defense d, the attack a is said to dominate attack a', denoted by $a \succeq^A_{p,u|d} a'$, if and only if

$$\int_{\Theta} u_A(d, a, \theta) p_A(\theta | d, a) d\theta > \int_{\Theta} u_A(d, a', \theta) p_A(\theta | d, a') d\theta, \ \forall (u_A, p_A) \in \mathcal{T}_A. \quad (1)$$

In particular, this dominance relation helps identify the attacks $a' \in \mathcal{A}$ that the Attacker will *not* choose in response to the defense $d \in \mathcal{D}$, because the Attacker, being a rational player, will not choose the attack a' if its expected utility for him is lower than that of $a \in \mathcal{A}$. Thus, partial information $\mathcal{T}_A$ about the Attacker type defines a strict partial order (an irreflexive and transitive binary relation) over the set of possible attacks.

3.1. Partial preference information and stochastic dominance

Using partial information to derive the dominance relation (1) is quite general and subsumes several cases of stochastic dominance for deriving a partial preference order over random variables (Levy, 1992). Specifically, assuming that $p_A = p$ so that the Attacker's probability estimates are known by the Defender, have consequences $c(d, a, \theta) \in \mathbb{R}$ and have the Attacker's utility function u_A belong to the set of all increasing utility functions $\mathcal{U}^+$, Eq. (1) becomes

$$\int_{\Theta} u_A(d, a, \theta) p(\theta | d, a) d\theta > \int_{\Theta} u_A(d, a', \theta) p(\theta | d, a') d\theta, \ \forall u_A \in \mathcal{U}^+.$$

In other words, every expected utility maximizing Attacker with an increasing utility function prefers attack a over a'. When attack alternatives are viewed as choices between random variables, this is equivalent to stating that a dominates a' in the sense of first degree stochastic dominance, denoted by $a \succeq_{\text{FSD}} a'$.

Similarly, if the Attacker is risk averse, then his utility function belongs to the set of all increasing concave utility functions and consequently and his decisions can be analyzed with second order stochastic dominance. Conversely, if the Attacker is risk prone, then $\mathcal{U}_A$ can be taken to be the set of all increasing convex utility functions. Different degrees of stochastic dominance can be used to describe how decision makers with different risk attitudes would rank decision alternatives resulting in uncertain outcomes.

If the consequences c involve multiple attributes c_i, it may not be known how important the different attributes are relative to each other from the Attacker's perspective. When $\mathcal{U}_A$ is the set of additive utility functions with increasing utilities for each attribute, aggregated with some non-negative (standardized) weights, the following Pareto-type first-order stochastic dominance holds

$$a \succeq_{\text{PFSD}} a' \iff p_A(c_i(d, a, \theta) \le e) \le p_A(c_i(d, a', \theta) \le e), \ \forall i, \forall e, \quad (2)$$

where the inequality is strict for some combination of consequences e and attributes i. While the assumption of additive multivariate utility functions does not always hold, the inaccuracies caused by its minor violations are often acceptable in practice (Keeney & von Winterfeldt, 2011) and can be rectified by restructuring the attributes. Note that Pareto dominance can be defined similarly for higher orders of stochastic dominance for real-valued attributes. It is also possible to construct more conclusive dominance relations, provided that partial information about the relative importance of the attributes can be obtained (See also: Liesiö, Mild, & Salo, 2008; Liesiö & Salo, 2012).

The construction of a stochastic dominance relation from a set of utility functions has been studied extensively; see, for example, Hadar and Russell (1969), Bawa (1975), Fishburn (1980), Kim (1998) and Shaked and Shanthikumar (2007). Thus, instead of focusing on the construction of dominance relations, we consider how they can be used in ARA.

3.2. Defense-Attack game with partial information about the Attacker

Using the basic sequential defend-attack game in Fig. 1 as an example, we analyse it as a game with partial information.

1. At the outcome node Θ, given decisions a and d, analyze the expected utility for the Attacker for each $(u_A, p_A) \in \mathcal{T}_A$
$$\psi_A^{u_A,p_A}(d,a) = \int_\Theta u_A(d,a,\theta)p_A(\theta|d,a)d\theta.$$
2. As the Attacker can observe the defense d chosen by the Defender, the set of non-dominated attacks $\mathcal{A}^*(d)$ at the node A can be computed based on the dominance relation $\succeq^A_{p,u|d}$ from (1) as
$$\mathcal{A}^*(d) = \{a : \nexists a' : a' \succeq^A_{p,u|d} a\}. \quad (3)$$
3. At node D, based on the Defender's expected utility
$$\psi_D(d,a) = \int_\Theta u_D(d,a,\theta)p_D(\theta|d,a)d\theta,$$
compute the set $\mathcal{D}^*$ of non-dominated actions in $\mathcal{D}$, using the dominance relation
$$d \succeq^D d' \iff \psi_D(d,a) \geq \psi_D(d',a'), \quad \forall a \in \mathcal{A}^*(d), \forall a' \in \mathcal{A}^*(d')$$
where the inequality is strict for at least one pair of attacks a, a'.

The key difference with the Bayesian analysis in Section 2.2 is that, rather than determining *how likely* the different actions of the Attacker are, the emphasis is on identifying *what* actions are plausible. The emphasis on this latter question is motivated by the fact that the Defender is ultimately interested not so much in what the Attacker will do as in maximizing her own expected utility.

If the Defender does not know the probability of the Attacker's responses $p_D(a|d)$, she cannot calculate her expected utility $\psi_D(d)$ exactly. However, the Defender can determine which of the Attacker's responses $\mathcal{A}^*(d)$ are possible so that they have a positive probability. The Defender can use this information to calculate an upper and a lower bound for her expected utility for each defense d, i.e.,

$$\psi_D^{\max}(d) = \max_{a \in \mathcal{A}^*(d)} \psi_D(d,a) \quad \text{and} \quad \psi_D^{\min}(d) = \min_{a \in \mathcal{A}^*(d)} \psi_D(d,a). \quad (4)$$

These upper and lower bounds can be used to check dominance relations between alternative defenses, because

$$d \succeq^D d' \iff \psi_D(d,a) \geq \psi_D(d',a'), \quad \forall a \in \mathcal{A}^*(d), \forall a' \in \mathcal{A}^*(d') \\ \iff \psi_D^{\min}(d) \geq \psi_D^{\max}(d').$$

This improves computational efficiency, because there is no need to compare the consequences for all possible responses $a \in \mathcal{A}^*(d)$ when determining $\mathcal{D}^*$. In some cases, these bounds can be determined analytically to further improve computational performance. If, for example, one of the non-dominated attacks a causes most harm to the Defender no matter what the Defender does (i.e., $\forall d \in \mathcal{D} : \exists a \in \mathcal{A}^*(d)$ such that $\psi_D(d,a) < \psi_D(d,a'), \forall a' \in \mathcal{A}^*(d)$), one need not compare all non-dominated attacks to establish the lower bound for ψ_D.

3.3. Partial information about both players

Sometimes, it may be impractical or impossible to specify the Defender's beliefs and preferences completely. This situation belongs to the realm of robust Bayesian analysis (Ríos Insua & Ruggeri, 2012). It can be analyzed by extending the results of the preceding section by assuming that

$$u_D(d,a,\theta) \in \mathcal{U}_D, \quad p_D(\theta|d,a) \in \mathcal{P}_D,$$

for the specified sets $\mathcal{U}_D, \mathcal{P}_D$ of possible utility functions and probability distributions that represent partial information about the Defender's preferences and beliefs, respectively. These sets define the set $\mathcal{T}_D \subseteq \mathcal{U}_D \times \mathcal{P}_D$ representing the Defender's possible types.

For comparing the Defender's defence alternatives, we define the dominance relation $\succeq^D_{p,u}$

$$(d,a)w \succeq^D_{p,u} (d',a') \iff \\ \int_\Theta u_D(d,a,\theta)p_D(\theta|d,a)d\theta \\ \geq \int_\Theta u_D(d',a',\theta)p_D(\theta|d',a')d\theta, \ \forall (u_D,p_D) \in \mathcal{T}_D, \quad (5)$$

where the inequality is strict for at least one combination $(u_D, p_D) \in \mathcal{T}_D$. For the Defender, this covers the cases of stochastic dominance in Section 3.2.

We now proceed as follows:

1. At the outcome node Θ, compute the Attacker's expected utility for
$$\psi_A^{u_A,p_A}(d,a) = \int_\Theta u_A(d,a,\theta)p_A(\theta|d,a)d\theta,$$
for each $(u_A, p_A) \in \mathcal{T}_A$ and $(a,d) \in \mathcal{A} \times \mathcal{D}$. and similarly for the Defender, for each feasible pair $(u_D, p_D) \in \mathcal{T}_D$.
2. At the Attacker's decision node A, compute the set of non-dominated attacks $\mathcal{A}^*(d)$ for the observed defense d.
3. At node D, compute the set $\mathcal{D}^*$ of non-dominated defenses in $\mathcal{D}$ based on the dominance relation
$$d \succeq^{D'} d' \iff (d,a) \succeq^D_{p,u} (d',a'), \quad \forall a \in \mathcal{A}^*(d), a' \in \mathcal{A}^*(d') \quad (6)$$
where at least one of the inequalities in (5) for the binary relation $(d,a) \succeq^D_{p,u} (d',a')$ is strict.

The approach in Section 3.2 cannot be applied here, because the upper and lower bounds for the Defender's expected utility cannot be calculated if the Defender's utility function is not known.

This notwithstanding, for some sets of utility functions an analogous approach may be taken. For example, if the consequences $c(d,\ a,\ \theta)$ are assessed with a single attribute and the Defender's preferences for these consequences are monotonic, then these preferences can be modeled with first-order stochastic dominance. It is also possible to establish the upper and lower bounds for the cumulative distribution functions

$$F_c^{\max}(c(\theta'|d)) = \max_{a \in \mathcal{A}^*(d)} \int_{-\infty}^{\theta'} p_D(c(d,a,\theta))d\theta \quad (7)$$

$$F_c^{\min}(c(\theta'|d)) = \min_{a \in \mathcal{A}^*(d)} \int_{-\infty}^{\theta'} p_D(c(d,a,\theta))d\theta.$$

and to use first-order stochastic dominance based on the comparison of bounds to determine the non-dominated decision alternatives

$$d \succeq^{D'} d' \iff (d,a) \succeq^{D}_{p,u} (d',a'), \forall a \in \mathcal{A}^*(d), \forall a' \in \mathcal{A}^*(d')$$
$$\iff F_c^{\min}(c(\theta'|d)) \le F_c^{\max}(c(\theta'|d)), \forall \theta',$$
$$\text{with strict inequality for some } \theta'.$$

Similar methods can also be applied to examine preferences for higher orders of stochastic dominance or Pareto dominance. In the first case, one can compute the upper and lower bounds for the integral that is used for the comparison, whereas in the additive multi-dimensional Pareto case it is necessary to establish bounds for the marginal cumulative distribution functions of each of the attributes.

3.4. Decision rules

The approaches in Sections 3.2 and 3.3 seek to provide as conclusive results as possible based on the available information. However, if the specification of the adversaries preferences and beliefs is very incomplete, the resulting sets of non-dominated defense and attack decisions may be too large to provide actionable decision support. This is a worthwhile result in and of itself, because it shows that the available information is insufficient for recommending a well-founded decision. If additional information about the players' types cannot be readily obtained, the set of recommended decisions can be narrowed down by introducing additional constraints. One can also apply decision criteria such as maximax, maximin or minimax regret from Bayesian robustness analysis (Ríos Insua & Ruggeri, 2012). See also the approach presented by McLay, Rothschild, and Guikema (2012), who consider consider robust optimization ideas in ARA contexts based on worst case scenarios.

For the Defender, the maximin decision rule for the set of non-dominated decisions is

$$\mathcal{D}^*_{\text{maximin}} = \max_{d\in\mathcal{D}^*} \min_{a\in\mathcal{A}^*(d)} \psi_D(d,a).$$

If the Defender's utility function is known, $\mathcal{D}^*_{\text{maximin}}$ contains a single alternative (or multiple equally preferred alternatives) which can be determined using the upper and lower bounds on the expected utility from Eq. (4)

$$\mathcal{D}^*_{\text{maximin}} = \max_{d\in\mathcal{D}^*} \psi_D^{\min}(d).$$

Optimal decision sets for maximax and minimax regret can be constructed similarly.

As in Section 3.3, the set $\mathcal{D}^*_{\text{maximin}}$ can be computed more efficiently if the Defender's preferences over the uncertain consequences fulfil first-order stochastic dominance. Then, the upper and lower bounds of the cumulative distribution functions from (7) can be employed to determine the non-dominated decision alternatives

$$\mathcal{D}^*_{\text{maximin}} = \{d \in \mathcal{D}^* : \nexists d' \in \mathcal{D}^* : F_c^{\max}(c(\theta|d)) > F_c^{\max}(c(\theta|d')), \forall \theta\}.$$

From the Defender's perspective, the upper bounds correspond to the worst possible combinations of the Attacker's and the Defender's preferences over consequences. The same approach can be applied to study other stochastic dominance relations.

In contrast to the use of decision rules in decision analysis, the Defender's choice between maximin, maximax and minimax regret decision rules may not reflect the Defender's risk attitude in the traditional sense as much as it reflects the Defender's aversion to the ambiguity associated with the Attacker's subsequent response. For example, the choice of the maximin decision rule limits the harm resulting from the Attacker's response to a minimum, whereas the minimax regret limits the downside variability in the Defender's expected utility. In some ARA problems, there could even be an ally instead of an Attacker, in which the Defender could choose the maximax decision rule to give the ally the opportunity to help her as much as possible.

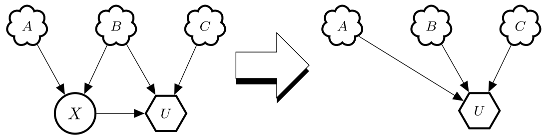

Fig. 2. Chance node removal.

3.5. Complex influence diagrams

Analyses based on partial information can be extended to more complex influence diagrams. In fact, the approach can be used to solve any *regular* multi-agent influence diagram in which the assumption of *no forgetting* holds. Regularity means that the influence diagram contains a directed path which traverses all decision nodes (regardless of which player they belong to) and thus defines a total order on them. The "no forgetting" assumption means that a player knows all the decisions and chance events that precede his/her current decision. Thus, such an influence diagram represents a sequential game of *perfect information*.

As shown by Shachter (1986) and Tatman and Shachter (1990), any regular influence diagram with no forgetting can be evaluated with a node elimination algorithm. Ortega, Rios, and Cano (2019) have extended this to bi-agent influence diagrams. A regular influence diagram can be solved with five graphical transformations:

1. Barren node elimination
2. Arc reversal between chance nodes
3. Chance node removal
4. Decision node removal
5. Value node removal

Barren node elimination removes non-utility nodes without children. The arc reversal transformation between chance nodes presented by Shachter (1986) is based on Bayes' rule and does not require any special considerations when dealing with partial information. The transformations for node removal, however, are more involved.

Fig. 2 shows the chance node removal operation. The cloud shaped nodes represent other parts of the influence diagram and can contain multiple decision and chance nodes as well as any number of arcs within themselves and between each other as long as the diagram remains regular. If a chance node has only children which are utility nodes belonging to different players, it can be removed by conditional expectation and the utility nodes inherit all the chance node's parents. In a multi-agent influence diagram, every utility node is associated with some player i with possible types $\mathcal{T}_i \subseteq \mathcal{U}_i \times \mathcal{P}_i$. As in (1), to remove the chance node X we compute a new partial preference order ranking the states of the new parent nodes A, B and C. That is, $(a,b,c) \succeq^i_{ABC} (a',b',c')$ if and only if

$$\int_{\mathcal{X}} u_i(b,c,x)p_i(x|a,b)dx \ge \int_{\mathcal{X}} u_i(b',c',x)p_i(x|a',b')dx, \ \forall (u_i,p_i) \in \mathcal{T}_i,$$

where the inequality is strict for at least one feasible pair of utility functions and probability distributions $(u_i,p_i) \in \mathcal{T}_i$. The partial preference order $\succeq^i_{ABC}$ is then used to form the new $\mathcal{T}'_i \subseteq \mathcal{U}'_i \times \mathcal{P}'_i$ for the new game represented by the modified influence diagram.

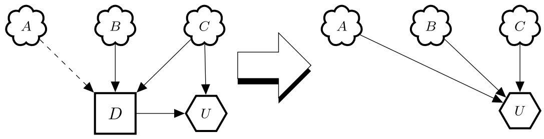

Fig. 3. Decision node removal.

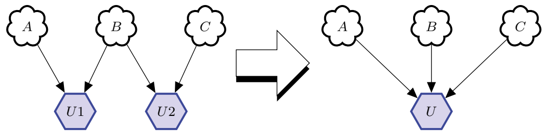

Fig. 4. Utility node removal.

This process is repeated for all the players whose utility nodes had X as a parent.

Fig. 3 illustrates the decision node removal operation. If all the children nodes for the decision node D for player i are utility nodes so that none of them have parents which are also shared by D, such decision node can be removed by maximization, in which case the utility node belonging player j inherits all the chance node's parents. (It is also possible that $i = j$.) As before, the players' possible types are represented by $\mathcal{T}_i \subseteq \mathcal{U}_i \times \mathcal{P}_i$ and $\mathcal{T}_j \subseteq \mathcal{U}_j \times \mathcal{P}_j$. We first determine non-dominated decision alternatives of player i at node D similarly as in Eq. (3)

$$\mathcal{D}_i^*(a, b, c) = \{d : \nexists d' : d' \succeq_D^i d\},$$

where $\succeq_D^i$ is derived from the utilities of player i. It is worth noting that if this player has no utility node as a child of D, then nothing is known about his decision and $\mathcal{D}_i^* = \mathcal{D}$. Utility node U is then updated based on $\mathcal{D}_i^*(a, b, c)$ similar to (6) in that

$$(a, b, c) \succeq_{ABC}^j (a', b', c') \iff$$
$$(d, c) \succeq_{DC}^j (d', c'), \forall d \in \mathcal{D}_i^*(a, b, c), \forall d' \in \mathcal{D}_i^*(a, b, c),$$

where $\succeq_{DC}^j$ represents the partial preference order of player j over states of C and D. It is reduced to an expected utility comparison if the type t_j for player j is known. The partial preference order $\succeq_{ABC}^j$ is then used to form the new $\mathcal{T}'_j \subseteq \mathcal{U}'_j \times \mathcal{P}'_j$ for the new game represented by the modified influence diagram. This process is repeated for all the players whose utility nodes have D as a parent.

Sometimes it may be convenient to represent player's utilities with several utility nodes. If the utility nodes share a parent node, it eventually becomes necessary to combine them to completely solve the influence diagram. Fig. 4 depicts how two utility nodes are combined into one and the new node inherits the parents of both. The partial preference relation describing the new utility node just has to be consistent with the removed ones

$$(a, b, c) \succeq_{ABC}^i (a', b', c') \iff$$
$$\left((a, b) \succeq_{AB}^i (a', b')\right) \wedge \left((b, c) \succeq_{BC}^i (b', c')\right).$$

Utility nodes belonging to different players cannot be normally combined in this way. Still, this does not prevent solving the influence diagram, because decision and chance nodes can be removed even if they have multiple different players' utility nodes as children.

With the last three graphical transformations now compatible with partial information, we can solve all sequential games of perfect information. The transformations can also help solve complex non-sequential games, but they cannot be used to eliminate decision nodes when an utility node has multiple decision node parents, as is the case in Fig. 5. All decisions that do not lie on the same directed path need to be solved simultaneously, and depending on the partial preference order at the utility node U and the number of decision nodes, this may be fairly straightforward or nearly impossible. For some preference orders it is possible to eliminate dominated pure strategies iteratively in order to reach an equilibrium solution (Börgers, 1994) For others, it is necessary to also consider mixed strategies (Perea, Peters, Schulteis, & Vermeulen, 2006). To our knowledge, non-sequential partial information games more complex than the one in Fig. 5 have not been studied.

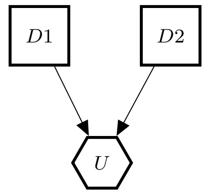

Fig. 5. Non-sequential decisions.

4. Planning of countermeasures for unmanned aerial vehicles

We illustrate our methods for solving complex influence diagrams with a realistic case study from military planning. We examine a scenario in which the Defender seeks to protect her supply company from the Attacker's reconnaissance activities. Specifically, the Attacker seeks to map the position of the company with UAVs. If the Attacker succeeds in this mapping activity, he can use either artillery or heavy rocket launchers against the Defender's supply company. The following ARA is produced for the Defender who seeks to assess the cost-efficiency of UAV-countermeasures for investment planning.

4.1. Scenario description

The scenario is shown as a bi-agent influence diagram in Fig. 6. The Defender has deployed a supply company around the village of Tarttila, Fig. 7. The Defender seeks to protect the company from artillery fire and hence also from UAV reconnaissance as cost-efficiently as possible. Towards this end, she considers two different anti-UAV weapon systems, two different radar systems, and the option of improving camouflage. She also has to choose how many weapon systems to buy and where to place them.

The Attacker seeks to destroy or cripple the supply company by inflicting losses through artillery fire. The Attacker knows only the general area where the company is located. The Attacker can employ three types of UAVs for reconnoitering artillery targets. All types of UAVs can be equipped with one of two different sensor systems. Because the UAVs cannot send information back in real time, they must survive long enough to return to their base. The UAVs cannot change their flight paths based on what they observe.

The Attacker has several artillery and rocket launcher systems with unguided ammunition. After examining the information provided by the UAVs, the Attacker decides how to employ these systems.

The Defender does not know how many UAVs or artillery systems the Attacker has, or how much it would cost for the Attacker to use or lose them. Nor is the Attacker's decision to use artillery or UAVs guided by these costs in the combat scenario. Overall, the Defender does not have complete information about the Attacker's utilities.

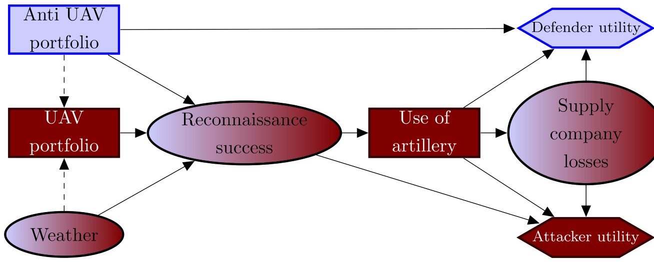

Fig. 6. Bi-agent influence diagram of the anti-UAV problem.

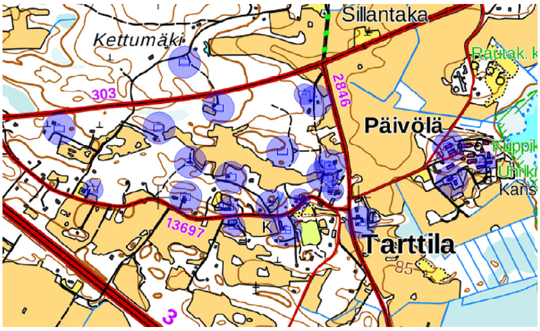

Fig. 7. The company in Tarttila. Blue circles indicate locations of platoons and squads. (Background map from of Finland, 2007).

4.2. Partial preference information

To keep this case study as realistic as possible, we use only as limited information about Attacker's preferences as one could expect to have in an actual military conflict. The Attacker wants to maximize the damage to Defender's equipment and personnel. At the same time he does not want to use ammunition unnecessarily for a given effect and prefers not losing UAVs. Because the Defender does not know what kind of UAV losses or ammunition costs the Attacker is willing to incur in order to cause to damage to the Defender, the Attacker's preferences over outcomes are formulated through Pareto-type first-order stochastic dominance (Eq. (2)). Specifically, the Attacker is assumed to (i) minimize ammunition consumption and UAV losses, and (ii) maximize the damage to the Defender's supply company. In addition, using Pareto dominance includes the assumption that the utilities for each of these attributes are additive and independent.

Because there can be a considerable delay between the Defender's investment into countermeasures and the actual deployment of these countermeasures, it can be challenging to formulate a utility function which accurately reflects trade-offs between immediate costs and uncertain future casualties. Thus, in the initial analysis, Pareto dominance is also employed to characterize the Defender's preferences for defense alternatives. Specifically, the Defender seeks to (i) minimize the losses to her supply company and (ii) maximize the amount of ammunition that the Attacker would have to use to cause such losses. The Defender also seeks to minimize the cost of her UAV countermeasures.

4.3. Artillery fire

Following the process described in Section 3.5, the first node to be eliminated from the influence diagram in Fig. 6 is *Supply company losses* from artillery fire. To eliminate the node, we calculated the conditional probabilities for equipment and personnel losses as a function of how the Attacker would aim his Artillery and how many shells or rockets he would fire. These probabilities were computed with the operation analysis software Sandis (Lappi, 2008).

We first chose 10 locations for the Defender's units that the Attacker might be able to identify using his reconnaissance. These same locations were used as targets for the artillery. It was determined the Attacker would be unlikely to accurately identify what equipment and personnel the Defender has at each of these locations, so all locations were deemed equally attractive to the Attacker. As can be seen from Fig. 8, the effects of the artillery fire increase when the Attacker uses more ammunition or divides fire between more targets. Unsurprisingly, damage to equipment and personnel were practically perfectly correlated. With this information, we can eliminate the chance node and update preferences for both sides. The preferences about supply company casualties now become irrelevant, because no other node affects them, so they are removed. The Attacker's preferences are updated and he wants to use as much ammunition as possible and spread it between as many locations as possible. The Defender, on the other hand, would prefer just the opposite.

The next node to be eliminated is the Attacker's decision node about how to *use artillery*. This decision depends on how many targets have been discovered by the UAV reconnaissance. The Attacker wants to use as much ammunition as possible to inflict maximum casualties, but, at the same time, he wants to conserve ammunition. These objectives are so obviously in conflict that no decision alternative is going to be dominated based on them. Because the Attacker prefers to spread artillery fire between as many targets as possible, he will always prefer to have as many as possible targets to aim at. Thus, when the Attacker's decision node is removed, he will no longer have preferences over aiming locations or ammunition consumption, but higher number of targets identified with UAV reconnaissance is now preferable. The Attacker also still wants to minimize UAV losses. The Defender wants to minimize the number of targets identified by the UAVs, maximize UAV losses, and minimize investment costs.

4.4. UAV reconnaissance

After eliminating the uncertain effects of the artillery fire and Attacker's decisions concerning it, the influence diagram has been reduced to a Defend-Attack game of UAV-reconnaissance, Fig. 9. The decision alternatives that remain for the Defender and Attacker at the remaining decision nodes are given in Tables 1 and 2, respectively. Because the Attacker can observe the Defender's decision, the Attacker's non-dominated responses to the Defender's decision are determined first.

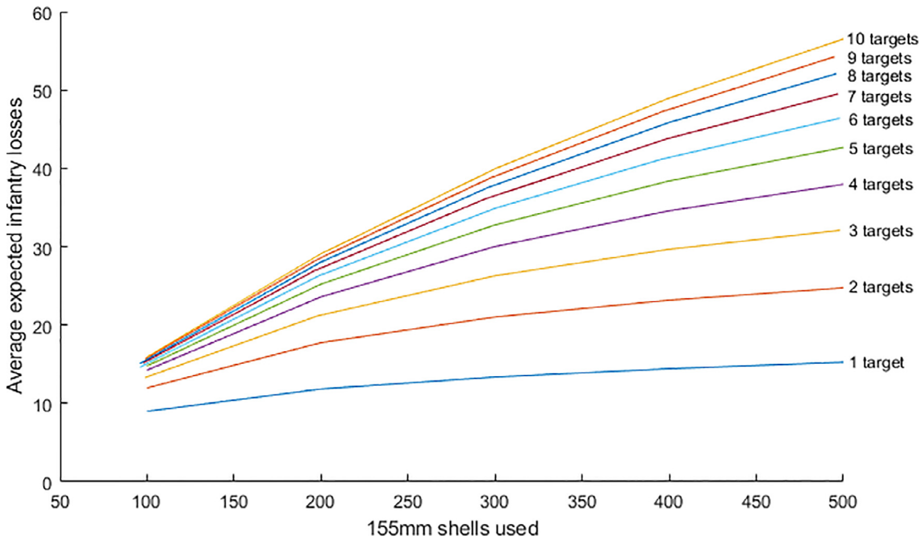

Fig. 8. Artillery fire inflicts higher total casualties when divided evenly between multiple targets.

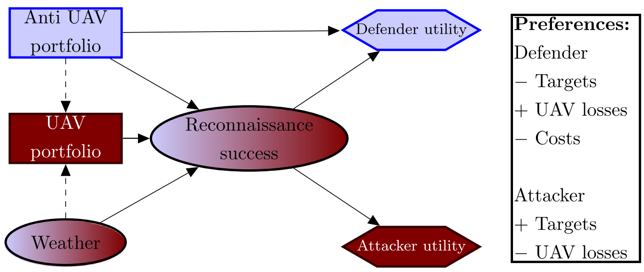

Fig. 9. Influence diagram after eliminating the artillery fire nodes.

Table 1
Defender's possible UAV countermeasures (see Appendix for technical details).

Decision	Description	Cost
Weapon system type	Projectile / Laser	110/160
Number of weapons	Up to 5	Per unit
Weapon locations	center, NW, NE, SW and SE	
Radar system type	Old radar / New radar	60/100
Number of radars	Up to 5	per unit
Radar locations	Center, NW, NE, SW and SE	
Added camouflage	Halves detection distance	130

Table 2
Attacker's UAV reconnaissance alternatives (see Appendix for technical details).

Decision	Description
UAV type	3 alternatives
Number of UAVs	Unlimited
Sensor type	2 alternatives
Flight altitude	Depends on UAV system

We built a MATLAB simulation tool to estimate the effectiveness of UAV reconnaissance. The tool incorporates a radar model that computes the detection probability based on physical and technical properties of the UAVs and radars and a simple weapon model that automatically destroys any UAVs that are within effective weapon range and have been detected. The conditional probability distributions for the detecting different numbers of target points were estimated through Monte Carlo simulation of random UAV flight paths.

The simulations were carried out for all combinations of Defender's and Attacker's decisions, making it possible to compute the Attacker's non-dominated response(s) to the Defender's decisions and weather conditions. The Attacker's decisions were evaluated based on the use of Pareto dominance for the two main objectives, i.e., the Attacker prefers to lose as few UAVs as possible and find as many target points as possible. No assumptions were made about how important these two objective would be to the Attacker.

The final step is to examine what decisions can be optimal in maximizing the Defender's expected utility. We did not elicit an exact utility function from the Defender, but we have a partial stochastic ordering based on the Defender's attribute specific preferences, similar to Section 3.3. Solving the problem using Pareto dominance over costs and the number of detected targets lead to 1271 non-dominated defense portfolios for the Defender. This is fewer than the 7689 at the outset, but still too many to recommend a decision. One reason why there are so many non-dominated portfolios is that the Attacker's response to the Defender's countermeasures is uncertain. This leads to wide bounds for the probabilities of finding different numbers of targets, see

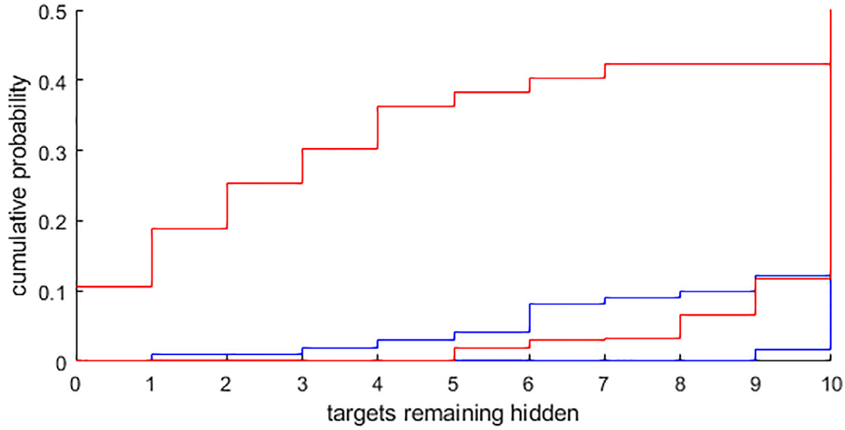

Fig. 10. Upper and lower bounds for two cumulative distribution functions (red and blue) for the number of target points that remain hidden for Defender's non-dominated alternatives. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3
Defender's non-dominated maximin decision alternatives.

Purchase cost	Weapon		Radar		Camouflage
	Model	Amount	Model	Amount	
0	–	0	–	0	No
170	Projectile	1	Old	1	No
210	Projectile	1	Old	1	Yes
230	Projectile	2	Old	1	No
290	Projectile	3	Old	1	No
300	Projectile	1	Old	2	No
320	Projectile	2	New	1	No
340	Projectile	1	New	1	Yes
420	Projectile	2	New	2	No
430	Projectile	3	New	1	No
450	Projectile	2	New	2	Yes

Fig. 10. Even though both the upper and lower bounds of the blue alternative shown in Fig. 10 are significantly lower than those corresponding to the red alternative, one cannot conclude which of these alternatives is better, because the Defender's decisions may affect those of the Attacker, and the upper bound for the blue alternative is above the lower bound for the red one.

To narrow down the set of plausible decisions, additional preference information from the Defender is needed. In this example, the Defender decides she wants to be prepared for worst-case situations, because the supply company needs to stay protected not only against a single attack, but also against repeated ones over a longer period. Thus, the maximin decision rule is applied to narrow down the non-dominated alternatives.

Because the Defender's utility function is not fully specified, the maximin decision rule does not yield a single decision. However, this rule reduces the number of non-dominated decision portfolios to 30 of which most differ only in the placement of the weapon systems. There are only 11 different equipment combinations, shown in Table 3. Many of the less costly alternatives, like the one which ignores UAVs entirely, do not protect the supply company well, but are non-dominated due to their low cost.

The results show that, for instance, the more expensive laser weapon system should not be chosen and the more expensive new radar is worthwhile only if the Defender is willing to spend at least 320 units. Because there are only a few alternatives, it would be possible to present the Defender with a detailed analysis of the probabilities with which the targets are discovered and what the likely effects of artillery fire are; or to present a well-founded overall assessment of what the most cost-efficient countermeasures are (Kangaspunta, Liesiö, & Salo, 2012).

Computing the non-dominated decision alternatives for the Defender took around 10 minutes using a fairly typical laptop and custom MATLAB code that was not optimized for speed. This is orders of magnitude less time than it took to compute the conditional probability tables for the chance nodes, which took days for the artillery fire and hours for the UAV reconnaissance using the same computer. Calculating the conditional probabilities for the effects of artillery fire turned out unnecessary in the end, because they only affected the decision dominance in very predictable ways. Using Bayesian Nash equilibrium to solve the same problem would have actually required those conditional probabilities.

For more complex problems, the computation time for this kind of dominance-based analysis increases much in the same way as in Bayesian Nash equilibrium analysis, i.e., it grows exponentially with the number of parents and children at decision nodes. However, the number of conditional probabilities in the chance nodes grows just as fast. This means that in many practical applications, the burden involved in eliciting them is likely to overshadow computational difficulties in solving the actual game.

5. Conclusions

ARA is a promising approach to the management of risks in application areas such as security and defense, because in contrast to standard game theoretic approaches, it does not necessitate strong assumptions about common knowledge. Still, modeling the adversaries' interlinked decision process can be challenging, given that in practice it may be exceedingly difficult or practically impossible to elicit complete information about the adversaries' preferences and beliefs.

Motivated by this, we have proposed dominance concepts and associated computational methods for characterizing and synthesizing partial information in ARA. Specifically, we have considered several variants of partial information which reflect different types of partial information in the context of the sequential Defend-Attack model. However, our approach can be readily extended to even to more extensive ARA models which involve sequential decision making, by following the principles discussed in the context of more complex influence diagrams. The methods presented are general enough that they can be used with either single or multi-attribute utilities.

We have also presented an illustrative case study on military planning in which the Defender seeks to protect its supply company from UAV surveillance. The salient elements of this case study are representative of the problems encountered in other ARA applications, for instance in the realm of cyber security. In particular, the proposed solution concepts for the analysis of partial information are likely to be useful when there are challenges in assessing the adversaries' multi-attribute utility functions and probability estimates. For instance, it would be difficult to predict how a cyber criminal would weigh the risks of getting caught against potential gains.

ARA models with partial preference information can also be used to address problems which involve multiple decision makers with different objectives. For instance, in environmental decision making there are often multiple stakeholders whose utility functions can be difficult to elicit completely (Hämäläinen, 2015). In such settings, a partial order for the decision alternatives could be built by eliciting and synthesizing partial information from these stakeholders.

The partial information approach may help simplify complex ARA problems even if the adversaries' preferences can be specified through utility functions. That is, because stochastic dominance does not require that the expected utilities are calculated exactly, there is no need to elicit complete probability information for the influence diagram either. As the case with artillery fire in the UAV case study shows, it may suffice to know the direction of change in the adversaries' utilities in response to changes in the probability parameters.

Analyses based on partial information can also be applied to problems involving sequential decisions by multiple actors. Because any inaccuracies in estimated utilities will propagate and accumulate in 'deep' influence diagrams containing long paths between decision and chance nodes, it may be advisable to err on the side of caution and produce initial analyses based on partial information. Then, if the partial information approach does not narrow down decisions sufficiently, one can revert back to the more traditional approach, elicit the utility functions and repeat the analysis using sets of non-dominated decisions. This will still require fewer probability estimates than the specification of full parameters for the original problem.

Acknowledgements

We thank the research project's advisory group, Esa Lappi, Markku Kujala, Sami Helle, Vesa RyynÃ, and Tero Majamaa, from the Finnish Defence Forces for their contribution to the development of the case study and for their expert advice. We also thank the Finnish Defence Research Agency, and Bernt Åkesson especially, for providing access to the operations analysis tool Sandis.

Appendix A. Technical parameters in UAV simulations

Table 4
UAV parameters.

Parameter	Small UAV	Cheap UAV	Fast UAV
Cost	Unknown	Unknown	Unknown
Radar cross-section	0.01 square meter	0.1 square meter	0.1 square meter
Speed	20 meter/second	30 meter/second	200 meter/second
Min altitude	20 meter	50 meter	50 meter
Max altitude	160 meter	500 meter	500 meter
Expensive sensor range	1000 meter	1000 meter	1000 meter
Cheap sensor range	500 meter	500 meter	500 meter

Table 5
Radar parameters.

Parameter	Old radar	New radar
Cost	60	100
Frequency	9.5 gigahertz	3.5 gigahertz
Peak power	80 Watts	60 Watts
Pulse duration	1 meter second	1 meter second
Net gain	15 decibels	15 decibels
Pulses integrated	1	10
Probability of false alarm	1E-6	1E-6
Elevation angle	30°	70°

Table 6
Weapon parameters.

Parameter	Projectile weapon	Laser weapon
Cost	110	160
Range	1000 meter	2000 meter
Limited by visibility	No	Yes

References

Antos, D., & Pfeffer, A. (2010). Representing Bayesian games without a common prior. In *Proceedings of the ninth international conference on autonomous agents and multiagent systems: volume 1-volume 1* (pp. 1457–1458). International Foundation for Autonomous Agents and Multiagent Systems.
Banks, D. L., Aliaga, J. M. R., & Ríos Insua, D. (2015). *Adversarial risk analysis*. Boca Raton, FL: CRC.
Bawa, V. S. (1975). Optimal rules for ordering uncertain prospects. *Journal of Financial Economics, 2*(1), 95–121.
Bier, V., Oliveros, S., & Samuelson, L. (2007). Choosing what to protect: Strategic defensive allocation against an unknown attacker. *Journal of Public Economic Theory, 9*(4), 563–587.
Börgers, T. (1994). Weak dominance and approximate common knowledge. *Journal of Economic Theory, 64*(1), 265–276.
Brown, G., Carlyle, M., Salmerón, J., & Wood, K. (2006). Defending critical infrastructure. *Interfaces, 36*(6), 530–544.
Cox, A. L. JR (2009). Game theory and risk analysis. *Risk Analysis, 29*(8), 1062–1068.
of Finland, N. L. S. (2007). Topographic map raster 1:50 000. https://tiedostopalvelu.maanmittauslaitos.fi/tp/kartta?lang=en. Accessed: 2017-6-16.
Fishburn, P. C. (1978). Non-cooperative stochastic dominance games. *International Journal of Game Theory, 7*(1), 51–61.
Fishburn, P. C. (1980). Stochastic dominance and moments of distributions. *Mathematics of Operations Research, 5*(1), 94–100.
Hadar, J., & Russell, W. R. (1969). Rules for ordering uncertain prospects. *The American Economic Review, 59*(1), 25–34.
Hämäläinen, R. P. (2015). Behavioural issues in environmental modelling–the missing perspective. *Environmental Modelling & Software, 73*, 244–253.
Harsanyi, J. C. (1967). Games with incomplete information played by Bayesian players, I–III part i. the basic model. *Management Science, 14*(3), 159–182.
Ríos Insua, D., & Ruggeri, F. (2012). *Robust Bayesian analysis*: vol. 152. New York: Springer Science & Business Media.
Kangaspunta, J., Liesiö, J., & Salo, A. (2012). Cost-efficiency analysis of weapon system portfolios. *European Journal of Operational Research, 223*(1), 264–275.
Keeney, R., & von Winterfeldt, D. (2011). A value model for evaluating homeland security decisions. *Risk Analysis, 31*(9), 1470–1487.
Kim, C. (1998). Stochastic dominance, Pareto optimality, and equilibrium asset pricing. *The Review of Economic Studies, 65*(2), 341–356.
Lappi, E. (2008). Sandis military operation analysis tool. In *Proceedings of the second nordic military analysis symposium, Stockholm, Sweden, November 17–18*.
Levy, H. (1992). Stochastic dominance and expected utility: Survey and analysis. *Management Science, 38*(4), 555–593.
Liesiö, J., Mild, P., & Salo, A. (2008). Robust portfolio modeling with incomplete cost information and project interdependencies. *European Journal of Operational Research, 190*(3), 679–695.
Liesiö, J., & Salo, A. (2012). Scenario-based portfolio selection of investment projects with incomplete probability and utility information. *European Journal of Operational Research, 217*(1), 162–172.
McLay, L., Rothschild, C., & Guikema, S. (2012). Robust adversarial risk analysis: A level-k approach. *Decision Analysis, 9*(1), 41–54.
Nikoofal, M. E., & Zhuang, J. (2015). On the value of exposure and secrecy of defense system: First-mover advantage vs. robustness. *European Journal of Operational Research, 246*(1), 320–330.
Ortega, J., Rios, D., & Cano, J. (2019). Bi-agent influence diagrams from an adversarial risk analysis perspective. *European Journal of Operational Research, 273*(3), 1085–1096.
Ozdaglar, A., & Menache, I. (2011). *Network games: theory, models, and dynamics*. Morgan & Claypool Publishers.

Perea, A., Peters, H., Schulteis, T., & Vermeulen, D. (2006). Stochastic dominance equilibria in two-person noncooperative games. *International Journal of Game Theory, 34*(4), 457–473.
Rass, S., König, S., & Schauer, S. (2017). Defending against advanced persistent threats using game-theory. *PLoS One, 12*(1), e0168675.
Rios, J., & Ríos Insua, D. (2012). Adversarial risk analysis for counterterrorism modeling. *Risk Analysis, 32*(5), 894–915.
Ríos Insua, D., Rios, J., & Banks, D. (2009). Adversarial risk analysis. *Journal of the American Statistical Association, 104*(486), 841–854.
Roponen, J., & Salo, A. (2015). Adversarial risk analysis for enhancing combat simulation models. *Journal of Military Studies, 6*(2), 82–103.
Rothschild, C., McLay, L., & Guikema, S. (2012). Adversarial risk analysis with incomplete information: A level-k approach. *Risk Analysis: An International Journal, 32*(7), 1219–1231.
Sevillano, J. C., Ríos Insua, D., & Rios, J. (2012). Adversarial risk analysis: The somali pirates case. *Decision Analysis, 9*(2), 86–95.
Shachter, R. D. (1986). Evaluating influence diagrams. *Operations research, 34*(6), 871–882.
Shaked, M., & Shanthikumar, J. G. (2007). *Stochastic orders*. New York: Springer Science & Business Media.
Tatman, J. A., & Shachter, R. D. (1990). Dynamic programming and influence diagrams. *IEEE Transactions on Systems, Man, and Cybernetics, 20*(2), 365–379.
Wang, S., & Banks, D. (2011). Network routing for insurgency: An adversarial risk analysis framework. *Naval Research Logistics (NRL), 58*(6), 595–607.
Washburn, A. R., Kress, M., et al. (2009). *Combat modeling*: vol. 139. New York: Springer.
Xu, J., & Zhuang, J. (2016). Modeling costly learning and counter-learning in a defender-attacker game with private defender information. *Annals of Operations Research, 236*(1), 271–289.
Zhuang, J., & Bier, V. M. (2007). Balancing terrorism and natural disasters defensive strategy with endogenous attacker effort. *Operations Research, 55*(5), 976–991.

Publication III

Ahti Salo, Edoardo Tosoni, Juho Roponen, and Derek W. Bunn. Using cross-impact analysis for probabilistic risk assessment. *Futures & Foresight Science*, 4(2), e2103, September 2021.

Received: 13 April 2021 | Revised: 26 August 2021 | Accepted: 27 August 2021
DOI: 10.1002/ffo2.103

ORIGINAL ARTICLE

WILEY

Using cross-impact analysis for probabilistic risk assessment

Ahti Salo[1] | Edoardo Tosoni[1] | Juho Roponen[1] | Derek W. Bunn[2]

[1]Department of Mathematics and Systems Analysis, Systems Analysis Laboratory, Aalto University School of Science, Aalto, Finland

[2]London Business School, London, UK

Correspondence
Ahti Salo, Department of Mathematics and Systems Analysis, Systems Analysis Laboratory, Aalto University School of Science, 00076 Aalto, Finland.
Email: ahti.salo@aalto.fi

Funding information
Finnish Research Programme on Nuclear Waste Management (KYT2022) 2019-2022

Abstract

Cross-impact analysis is widely employed to inform management and policy decisions based on the formulation of scenarios, defined as combinations of outcomes of relevant uncertainty factors. In this paper, we argue that the use of nonprobabilistic variants of cross-impact analysis is problematic in the context of risk assessment where the usual aim is to produce conservative risk estimates which may exceed but are not smaller than the actual risk level. Then, building on the characterization of probabilistic dependencies, we develop an approach to probabilistic cross-impact analysis which (i) admits several kinds of probabilistic statements about the outcomes of relevant uncertainty factors and their dependencies; (ii) maps such statements into constraints on the joint probability distribution over all possible scenarios; (iii) provides support for preserving the consistency of elicited statements; and (iv) uses mathematical optimization to compute lower and upper bounds on the overall risk level. This approach—which is illustrated with an example from the context of nuclear waste repositories—is useful in that it retains the informativeness of cross-impact statements while ensuring that these statements are interpreted within the coherent framework of probability theory.

KEYWORDS
cross-impact analysis, probabilistic risk assessment, scenario analysis

1 | INTRODUCTION

In its many variants, scenario analysis is widely employed to support strategic decisions whose impacts depend on key uncertainties (Bunn & Salo, 1993; Lord et al., 2016). In such situations, the systematic identification of relevant uncertainty factors; the characterization of outcomes which depict possible realizations of these factors; and the formulation of scenarios as different combinations of such outcomes provides support for organizational learning, fosters managerial insights and provides an improved basis for strategic decisions through a systematic analysis of uncertainties (Schoemaker, 1993; A. Wright, 2005).

Yet, a practical challenge in scenario analysis is that the number of possible scenarios grows very rapidly with the number of uncertainty factors and their outcomes. This is because for every combination of outcomes of these uncertainty factors, there exists a distinct scenario that could be generated (Carlsen et al., 2016; Tietje, 2005). Thus, if there are 10 factors with five possible outcomes for each, for example, the total number of possible scenarios which can be defined by such outcome combinations is $5^{10} \approx 9.7$ million. Understandably, the number of scenarios which are usually elaborated is typically much smaller, given that resources for developing scenarios by engaging experts or by consulting other sources of information are limited. Moreover, the elaboration of scenarios and the assimilation of their implications is constrained by the amount of time and attention that decision and policy makers can devote to the scenario process. Thus, in many public policy and corporate scenario analyses which are developed primarily by consulting experts and other respondents, the number of scenarios is in the range between four and eight (see, e.g., Lord et al., 2016; Wiebe et al., 2018).

In this setting, methods of cross-impact analysis provide a structured approach to choosing those outcome combinations for which scenarios are built, based on statements concerning the logical relationships between the factors and their outcomes. Such statements are typically elicited by asking the respondent to characterize which pairs of outcomes are consistent in the sense that these outcomes are likely to occur jointly. Typically, these cross-impact statements are expressed verbally and then mapped to corresponding numerical parameters. For instance, Scholz and Tietje (2001) present a 7-point numerical scale from −3 to 3 such that, for example, −3 indicates that the two outcomes are *strongly inconsistent* in the sense that they are very unlikely to occur together; 0 represents *independence*; and 3 indicates that the outcomes are *strongly consistent* so that the occurrence of an outcome induces the other. Finally, the elicited statements are synthesized algorithmically to provide suggestions for which combinations of outcomes scenarios should be built (see, e.g., Salo & Bunn, 1995; Seeve & Vilkkumaa, 2021; Tietje, 2005).

As one of the important application areas of scenario analysis, risk assessment covers both risk analysis (which helps identify, characterize, and analyze future events and developments that can negatively impact individuals, assets or the environment) and risk evaluation (which supports judgments about the extent to which these risks can be tolerated) (Rausand, 2013). In risk assessment, the demands on the rigor, quality and transparency of methodological support are particularly stringent. In part, this is because risk management decisions can have far-reaching consequences, especially in the context of safety-critical systems whose failures can cause human casualties, irreversible environmental damages, and major financial losses. Thus, for example, in the assessment of the safety of nuclear waste repositories, it is necessary to account for the full range of relevant uncertainty factors (called features, events, and processes [FEPs]; see Tosoni et al., 2018) and their implications for regulatory decisions. Methodological rigor is also needed in assessing risks due to the impacts of climate change, healthcare interventions, and environmental regulations (see, e.g., Hirabayashi et al., 2013). In all these areas, the possibility of rare but extremely serious events is of much concern. These events have usually very low probabilities which can be notoriously difficult to estimate because of scarce empirical evidence and paucity of relevant data based (see, e.g., Goodwin & Wright, 2010).

Within the field of risk assessment, probabilistic risk analysis (PRA) constitutes a theoretically coherent framework which is compatible with well-established statistical techniques for data analysis; it also provides support for synthesizing expert judgments (Bedford & Cooke, 2001). In the analysis of safety-critical systems, it is often required that the PRA estimates—which reflect both the *probability* and the *severity* of negative impacts—should be *conservative* so that the actual risk level is not underestimated (see, e.g., Aven & Zio, 2011). This requirement is justified by the recognition that in safety-critical systems, errors due to “false negatives”—the failure to take appropriate risk management decisions in response to risks which were deemed tolerable but were actually too high—can be far greater than errors arising from “false positives”—the cost of unnecessarily implementing risk management actions in response to assessed risks which, in reality, were not big enough to warrant such actions. Even more generally, such conservatism is widely called for in situations where there are significant uncertainties. For example, the “precautionary principle” (Science for Environment Policy, 2017) has been invoked to guide the public response to risks in contexts such as climate change mitigation (Stern, 2007). Also the “minimax regret” decision rule, which has been proposed as an approach for ensuring the resource adequacy of electricity systems (National Grid, 2020), is motivated by the desire to limit the amount of harm that could be experienced ex post. If the impacts can be characterized in terms of real-valued consequences (for instance through monetization), information about the tail risk represented by the least preferred consequences can be provided through risk measures such as Value-at-Risk and conditional Value-at-Risk, defined at appropriate confidence levels (see Liesiö & Salo, 2012).

The above remarks motivate our central observation on the use of cross-impact analysis in risk assessment and the ensuing decision making. That is, to the extent that cross-impact analysis focuses on a small subset of all possible scenarios, there is a real possibility that the resulting estimates about the overall risk level will *not* be conservative, because the risks associated with all the other “non-constructed” scenarios may be underestimated or even neglected. This may not be of major concern in contexts where the stakes are not very high or where “softer” process objectives such as organizational learning are dominant. However, if the analysis serves as an essential input to safety-critical risk management decisions, it is possible that the sufficient conservatism required by regulatory decision making is not being upheld. Indeed, while all model-based analyses are simplifications and there is always some “model risk,” in safety-critical applications, due-diligence requires that this should be minimized.

Against this backdrop, we examine cross-impact analysis from the PRA perspective, with the aim of clarifying how cross-impact analysis can be employed to support risk management decisions. This perspective is motivated by the recognition that (i) risk assessment is, by definition, focused on the identification, characterization, and analysis of relevant uncertainties and their impacts, and that (ii) PRA is often endorsed and in many cases even required as the only appropriate coherent framework for addressing these uncertainties (see, e.g., Helton & Sallaberry, 2009; USEPA, 2014; USNRC, 2016). As a motivating prelude to our methodological development, we point out limitations in nonprobabilistic cross-impact approaches by examining the cross-impact balances (CIB) method (Weimer-Jehle, 2006, 2008). We have chosen this method due its visibility in the literature and the attention that it has recently received in the context of climate change mitigation (Kemp-Benedict et al., 2010; Panula-Ontto et al., 2018; Schweizer, 2020; Weimer-Jehle et al., 2020).

Furthermore, by building on formulations for capturing probabilistic dependencies, we develop a probabilistic method of cross-impact analysis which combines methodological coherence with the expressiveness of cross-impact statements for characterizing dependencies between pairs of outcomes for uncertainty factors. These statements are translated into constraints on the joint probability distribution over the set of *all* possible scenarios (which, by design, are assumed to be mutually exclusive and collectively exhaustive; see, e.g., the early work of Duperrin & Godet, 1975 and citations to it). In addition to cross-impact statements, our method accommodates many other kinds of probabilistic statements, such as lower or upper bounds on the marginal and conditional probabilities of the joint probability distribution. Throughout the elicitation process, the method can offer support for preserving the consistency of the elicited statements so that the corresponding constraints are satisfied by at least some scenario probabilities.

In the context of risk assessment, our method can also be employed together with measures of risk importance to identify the scenarios which matter most from the risk management perspective (see, e.g., Salo et al., 2021; Tosoni, 2021). A precondition for this is that estimates about the expected consequences in every possible scenario can be assessed. While the generation of such estimates can be supported by computational models in some contexts (cf. the case study in Section 4), this assessment task may be challenging if the number of possible scenarios is large and the required estimates have to be elicited from experts (see, e.g., Dias et al., 2018). This task may be less onerous if the consequences depend primarily on few uncertainty factors, because it may suffice to assess consequences by conditioning these on, say, pairs or triplets of outcomes for two or three uncertainty factors. It may also be possible to estimate scenario-specific consequences by using mathematical models in which the consequences are expressed as functions of the outcomes that define the scenarios. One possibility is to apply the rank nodes method (Fenton et al., 2016; Laitila & Virtanen, 2016) which has been successfully employed to support the development of conditional probability tables for Bayesian networks. This method appears particularly relevant thanks to its flexibility which is achieved by associating weighting parameters with each uncertainty factor.

More generally, even if scenario-specific consequences are not formally assessed, the proposed approach to the elicitation of cross-impact statements and their conversion into constraints on the underlying joint probability distribution provides a structured and systematic way for characterizing this distribution. In this regard, it serves similar purposes as approaches for modeling dependencies between continuous random variables with real-valued outcomes (see, e.g., Van Dorp, 2005).

While our emphasis is on probabilistic approaches, we note that nonprobabilistic approaches such as CIB do not automatically lead to excessively permissive conclusions about system safety, provided that deliberate attempts are made to select those scenarios which pose significant risks while also accounting for the impacts of those scenarios which are not elaborated. This notwithstanding, a major shortcoming of these nonprobabilistic approaches is that they are not founded on a coherent theoretical framework within which the adequacy, appropriateness, and sufficiency of these kinds of adjustments could be formally assessed. This makes it hard if not impossible to ascertain if such adjustments warrant valid conclusions about system safety. Thus, there is a striking contrast with PRA which, due to its probabilistic foundations, builds on a coherent framework within which such an assessment can be made.

The rest of this paper is structured as follows. Section 2 discusses methods of cross-impact analysis and remarks on nonprobabilistic approches in light of the CIB method. Section 3 shows how cross-impact statements can be converted into constraints on the joint probability distribution over all possible scenarios. It also formulates maximization problems which can be solved to infer conservative risk estimates, based on all the elicited information. Section 4 presents a numerical example. Section 5 concludes.

2 | METHODS OF SCENARIO AND CROSS-IMPACT ANALYSIS

Of the variety of methods in scenario analysis, most are associated with one of the three main schools which are commonly referred to as the intuitive logics school; the probabilistic/modified trends school; and La Prospective (Bradfield et al., 2005; Bunn & Salo, 1993). The first, intuitive logics, is least quantitative in that it adopts a top-down inductive approach in seeking to formulate descriptive scenarios which represent possible futures and thus help generate actionable insights (Bowman, 2016; G. Wright et al., 2013). The second school consists of methods such as Trend-Impact Analysis and Cross-Impact Analysis which employ techniques for quantifying expert judgments, for example by characterizing possible deviations from historical averages or prior expectations (Bradfield et al., 2005). The third school, La Prospective, can be viewed as a “blend of tools and systems analysis” (Godet, 2000) or even as a mixture of methods from intuitive logics and probabilistic analysis (Bradfield et al., 2005).

Regardless of the school, it is useful to consider the determinants of the scenario quality (Bunn & Salo, 1993). In particular, scenarios should be *comprehensive*, meaning that they represent the full range of possible futures that are relevant to decision making or the broader objectives of the scenario process; *consistent*, meaning that the outcome combinations are plausible in light of available knowledge about the reality which they seek to depict; and *coherent*, meaning that the development of scenarios is founded on sound theories for reasoning about uncertainties. In practice, the pursuit of these qualities involves inevitable trade-offs. For example, increasing the number of scenarios to ensure comprehensiveness would, at some stage, result in the generation of scenarios which are less plausible and therefore less consistent, too.

In our methodological development, we focus on probabilistic approaches in which uncertainty factors are modeled as random variables $X^i, i = 1, \dots, n$ such that the ith uncertainty factor has n_i possible realizations (called outcomes) $x^i_k, k = 1, \dots, n_i$ represented by the set $S^i = \{x^i_1, \dots, x^i_{n_i}\}$. A *scenario* $s = (x^1, \dots, x^n)$ is defined as a combination of outcomes $x^i \in S^i$ for all uncertainty factors $i = 1, \dots, n$. Thus, mathematically, the set of all scenarios is the Cartesian product $S = \bigtimes_{i=1}^{n} S^i$ which has $|S| = \prod_{i=1}^{n} n_i$ elements. For example, if there are 5 factors with three outcomes for each, there are $3^5 = 243$ distinct scenarios that can be generated.

Much of the early methodological development of cross-impact analysis took place in the 1970s and 1980s. One of the major aims was to support inferences about which scenarios could be deemed more plausible than others, based on cross-impact statements about the consistency of outcomes for pairs of uncertainty factors. The proposed methods were largely developed within the framework of probability theory by interpreting the elicited cross-impact judgments in terms of statements about conditional probabilities and by translating these statements into corresponding constraints on the joint probability distribution over the set of all possible scenarios (for an early review, see tab. 1 in Salo & Bunn, 1995).

In the elicitation of cross-impact statements, one notable challenge is that when several cross-impact statements are elicited without explicit guidance, the resulting set of elicited statements may be inconsistent so that the corresponding constraints will not be satisfied by any probability distribution over the set of scenarios (G. Wright et al., 1988). In this case, it would be necessary to revise earlier statements, either by removing some of them or, alternatively, by relaxing the bounds of those statements which have been encoded as intervals. Both cases are problematic in that it can be challenging to identify which one(s) of the many earlier statements are more "wrong" than others.

In recent years, the literature on cross-impact analysis has continued to diversify. There are now approaches in which the assessed cross-impact evaluations are no longer linked to probabilities. One of these approaches is the CIB method (Weimer-Jehle, 2006, 2008) which is a structured technique for identifying consistent scenarios based on cross-impact assessments about causal dependencies between uncertainty factors. In CIB, specifically, the respondent is invited to use a scale ranging from −3 to 3 to assess what impact the outcome $x_k^i \in S^i$ of the ith factor will have on the outcome $x_l^j \in S^j$ of the jth factor. These statements are assessed for all pairs of outcomes $(x_k^i, x_l^j), x_k^i \in S^i, x_l^j \in S^j$ and pairs of uncertainty factors $i \neq j$, resulting in responses $C_{kl}^{ij}, i \neq j, k = 1, \ldots, n_i, l = 1, \ldots, n_j$. These responses form the elements of the cross-impact matrix C.

In the selection of scenarios, CIB focuses exclusively on consistent scenarios which are defined as combinations of outcomes $\left(x_{k_1^*}^1, \ldots, x_{k_n^*}^n\right)$ such that (see eq. 1 in Weimer-Jehle, 2008)

$$\sum_{\substack{i=1 \\ i \neq j}}^{n} C_{k_i^* k_j^*}^{ij} \geq \sum_{\substack{i=1 \\ i \neq j}}^{n} C_{k_i^* l}^{ij}, j = 1, \ldots, n, l = 1, \ldots, n_j. \tag{1}$$

In other words, the scenario $\left(x_{k_1^*}^1, \ldots, x_{k_n^*}^n\right), x_{k_i^*}^i \in S^i$ is consistent in the sense that the sum of corresponding cross-impact terms in each column $x_{k_j^*}^j, j = 1, \ldots, n_j$ of the aggregate matrix is not less than what would be obtained by adding the terms in the column for some other outcome $x_l^j \neq x_{k_j^*}^j$ instead.

Even if this requirement seems plausible, it is highly restrictive in that the number of scenarios which satisfies the condition (1) can be very small, which undermines the objective of generating a *comprehensive* set of scenarios. For instance, in the example in tab. 3 of Weimer-Jehle (2006) with five factors (four with three possible outcomes and one with four), only three out of the $3^4 \times 4 = 324$ scenarios are consistent, because none of the 321 other scenarios fulfill the consistency requirement (1).

Alarmingly, it is also possible to construct cross-impact matrices such that the consistency requirement in (1) is not satisfied by *any* scenario. For example, consider the cross-impact matrix in Figure 1 which is based on two uncertainty factors such that the possible outcomes of the first factor are $\{a, b, c\}$ and those of the second factors are $\{x, y, z\}$. Then, condition (1) means that for example, the scenario (k_1^*, k_2^*) would be consistent if and only if $C_{k_2^* k_1^*}^{21} \geq C_{k_2^* l}^{21}, l \neq k_1^*$ for $j = 1$ in (1), and $C_{k_1^* k_2^*}^{12} \geq C_{k_1^* l}^{12}, l \neq k_2^*$ for $j = 2$ in (1).

Yet the following nine inequalities show that for *any* scenario there exists some other column such that at least one of these conditions is violated:

$$\begin{aligned}
&C_{ax}^{12} = 0 < 1 = C_{ay}^{12}, C_{ya}^{21} = -3 < 3 = C_{yc}^{21}, C_{az}^{12} = -1 < 1 = C_{ay}^{12} \\
&C_{bx}^{12} = -3 < 3 = C_{bz}^{12}, C_{by}^{12} = 0 < 3 = C_{bz}^{12}, C_{zb}^{21} = 0 < 1 = C_{za}^{21} \\
&C_{xc}^{21} = -1 < 1 = C_{xb}^{21}, C_{cy}^{12} = 0 < 1 = C_{cx}^{12}, C_{cz}^{12} = -2 < 1 = C_{cx}^{12}.
\end{aligned}$$

Even if the numerical values in the cross-impact matrix in Figure 1 are hypothetical, this example shows that there can be data sets of cross-impacts statements such that no scenarios satisfy the condition (1). Admittedly, the absence of consistent scenarios can be attributed to the lack of consistency in the statements. However, to the extent that the elicitation process offers no structured guidance for the specification of statements, there is a risk that the set of scenarios which are screened for further elaboration becomes too small, thus undermining the attainment of the comprehensiveness as a quality attribute. In other words, the strong emphasis on the consistency criterion based on a dichotomous "yes-no" assessment may, depending on the elicited cross-impact statements, be so stringent that the number of consistent scenarios is too small to ensure the comprehensiveness of the generated scenarios, all the more so because the extent to which the scenarios are comprehensive is not formally defined. From this perspective, we find that among nonprobabilistic

		Factor 1			Factor 2		
		a	*b*	*c*	*x*	*y*	*z*
Factor 1	*a*				0	1	-1
	b				-3	0	3
	c				1	0	-2
Factor 2	*x*	0	1	-1			
	y	-3	0	3			
	z	1	0	-2			

FIGURE 1 An example of inconsistencies

cross-impact methods there are significant advantages to adopting approaches which (i) employ quantitative measures for concepts such as consistency and comprehensiveness and (ii) provide suggestions for the selection of scenarios by solving corresponding optimization problems. One such approach for generating scenarios which are both plausible and diverse is presented in (Seeve & Vilkkumaa, 2021).

In Figure 1, the cross-impact terms are not monotonic in the sense that transitions to a higher index (e.g., moving first from a to b and then proceeding to c) would be associated with systematic increases or decreases in the assessed cross-impacts. In effect, the monotonicity of such changes makes sense only on condition that there exists a corresponding metric or ordinal scale such that there is a sense of direction ranging from outcomes on the lower levels to those on the higher levels (as opposed to a nominal scale which merely indicates selections from the set of outcomes without such directionality, for example, choices among political parties; see Carlsen et al., 2016).

Uncertainty factors which are assessed using metric scales (e.g., temperature) can be discretized to formulate corresponding ordinal scales. Then, assuming that there are such ordinal scales for all uncertainty factors, the outcomes for each factor can be ordered with a transitive, antisymmetric, and total binary relation $\prec_i, i = 1, \ldots, n$. In this case, the monotonicity property can be stated as

$$x^i_k \prec_i x^i_{k'} \prec_i x^i_{k''} \Rightarrow \left(\left[C^{ij}_{kl} \le C^{ij}_{k\&\#x00027;l} \le C^{ij}_{k''l}\right] \vee \left[C^{ij}_{kl} \ge C^{ij}_{k\&\#x00027;l} \ge C^{ij}_{k''l}\right)\right) \wedge \left(\left[C^{ji}_{lk} \le C^{ji}_{lk'} \le C^{ji}_{lk''}\right] \vee \left[C^{ji}_{lk} \ge C^{ji}_{lk'} \ge C^{ji}_{lk''}\right]\right). \quad (2)$$

In risk assessment, one should be wary of assuming that the lowest and highest risks would be attained at the endpoints of any such ordinal scale. For example, if departures from the normal conditions in a production facility are measured on a natural ordinal scale (or even an interval scale, as in the case of, e.g., temperature), deviations into either direction can contribute to increased risks.

Still, even with monotonic cross-impacts, it is possible that there are no CIB-consistent scenarios. One such example is in Figure 2 where there are three uncertainty factors whose outcomes belong to the sets $\{a, b, c\}$, $\{i, j, k\}$, and $\{x, y, z\}$, respectively. The shaded rows indicate the selection of the scenario (a, i, x) which is also indicated by the upward pointing arrows and the digits "1" in the second row at the bottom of the figure. The numbers in the first row under the downward arrows show the sums for those columns which have the highest column sum for the selection of outcomes for each uncertainty factor. For factor 3, this sum is the highest $2 = 0 + 2$ (obtained from matrix entries $C^{13}_{az} = 0$ and $C^{23}_{iz} = 2$) while the corresponding sum associated with the scenario (a, i, x) is $-4 = -3 + (-1)$, based on $C^{13}_{ax} = -3$ and $C^{23}_{ix} = -1$. Thus, scenario (a, i, x) is not consistent, because condition (1) would be violated by replacing the outcome x by z. It straightforward to check that none of the 27 scenarios are consistent.

In view of these examples, the procedures of the CIB method seem excessively restrictive in that there are examples of numerical inputs such that the consistency requirements hold either for very few or, at the limit, no scenarios at all. As a result, it appears that in the case of nonprobabilistic cross-impact analysis, approaches which are based on the formulation of optimization problems towards the identification of a set of consistent and diverse scenarios should be preferred. For example Seeve and Vilkkumaa (2021) present a structured approach which was applied to generate scenarios for the National Emergency Supply Agency in Finland. In what follows, however, we explore how the probabilistic interpretation of cross-impact statements can be employed to establish a coherent methodological foundation for using cross-impact analysis in the context of risk assessment, in particular.

2.1 | Probabilistic dependencies

There is an extensive literature on the characterization of probabilistic dependencies between events. Such dependencies will arise if there are causal relationships between the events; but they may very well exist even in the absence of such relationships. Specifically, research on the topic of probabilistic causation has sought to characterize what causation means in probabilistic terms (see, e.g., Williamson, 2009 for an overview as well as contributions by Pearl, 2013; Suppes, 1970). In general, there is wide agreement that a

		Factor 1			Factor 2			Factor 3		
		a	b	c	i	j	k	x	y	z
Factor 1	a				1	-3	-3	-3	0	0
	b				2	0	-2	3	0	-1
	c				3	2	1	3	-2	-2
Factor 2	i	3	3	-2				-1	0	2
	j	0	0	3				-1	0	1
	k	-3	0	3				-3	-2	-1
Factor 3	x	3	2	-2	1	0	0			
	y	3	0	-2	3	0	0			
	z	-3	-3	3	3	3	0			
		↓			↓					↓
		6	5	-4	2	-3	-3	-4	0	2
		1	0	0	1	0	0	1	0	0
		↑			↑			↑		

FIGURE 2 An example of inconsistencies with monotonic cross-impacts

statement such as "the event A may have been caused by the event B" can be interpreted as meaning that the occurrence of A is more likely if the event B has occurred, that is,

$$\mathbb{P}(A|B) > \mathbb{P}(A). \tag{3}$$

Here, the qualification "may have been" is warranted, because the inequality (3) lacks any contextual knowledge. For instance, it does not consider *when* the events occur, even if the attribution of causality would be possible only on condition that the event B occurs before A. Moreover, even if this were to be the case, it could be that the event A can be more meaningfully attributed to intermediate events which occur after B but before A. There are even parallels to empirical econometrics where the notion of "Granger causality" (Granger, 1969) is defined so that B is said to cause A if the regression $A(t) = a + bB(t-1)$ (where t refers to points in time) has a significant regression coefficient b but $B(t) = a + bA(t-1)$ does not.

In view of (3), we interpret the ratio $\mathbb{P}(A|B)/\mathbb{P}(A)$ as an indication of the degree of probabilistic dependency between the occurrence of events B and A, noting that this ratio need not be reflect causal relationships between the events. In keeping with this interpretation, we suggest that the cross-impacts are linked to ratios between conditional and marginal probabilities, as defined by

$$C^{ij}_{kl} := \frac{p^{i|j}_{k|l}}{p^i_k} = \frac{p^{ij}_{kl}}{p^i_k p^j_l}, \tag{4}$$

where $p^i_k = \mathbb{P}(X^i = x^i_k)$, $p^j_l = \mathbb{P}(X^j = x^j_l)$, $p^{i|j}_{k|l} = \mathbb{P}(X^i = x^i_k | X^j = x^j_l)$, and $p^{ij}_{kl} = \mathbb{P}(X^i = x^i_k \wedge X^j = x^j_l)$. In particular, C^{ij}_{kl} thus provides an answer to the question "How many times more likely does the outcome x^i_k of the ith uncertainty factor become if it is known that the outcome of the jth uncertainty factor is x^j_l?" This question invites intuitively meaningful and theoretically well-defined answers on a ratio scale. Such answers can be encoded with the help of verbal descriptors that can be calibrated through experiments (see Pöyhönen et al., 1997). Note that if the outcomes x^i_j are x^j_l are independent, then $p^{i|j}_{k|l} = p^i_k$ and $C^{ij}_{kl} = 1$.

Based on the interpretation (3), the cross-impact terms are symmetric, because (4) implies $C^{ij}_{kl} = C^{ji}_{lk}$. This property is desirable in that symmetry is aligned with the nondirectional relational structure of (in) consistencies. That is, stating that the events A and B are "inconsistent" does not involve causal judgments about why the joint occurrence is very unlikely or, in particular, whether or not it is the occurrence of one which is preventing the other from occurring. Furthermore, this property also makes it easier to elicit the cross-impacts terms, because evaluations are needed only for *unordered pairs* of outcomes (i.e., $\sum^n_{i,j=1,i\neq j}(n_i \times n_j)/2$) instead for all ordered pairs (i.e., $\sum^n_{i,j=1,i\neq j} n_i \times n_j$).

The following result shows that the relation (3) implies $\mathbb{P}(A|B) > \mathbb{P}(A|\neg B)$ and vice versa.

Theorem 1. *Assume that events A, B are such that* $0 < \mathbb{P}(B) < 1$. *Then*

$$\mathbb{P}(A|B) > \mathbb{P}(A) \Leftrightarrow \mathbb{P}(A|B) > \mathbb{P}(A|\neg B). \tag{5}$$

Proof. "$\Rightarrow$": If (3) holds, then

$$\mathbb{P}(A) = \mathbb{P}(A|B)\mathbb{P}(B) + \mathbb{P}(A|\neg B)\mathbb{P}(\neg B) > \mathbb{P}(A)\mathbb{P}(B) + \mathbb{P}(A|\neg B)\mathbb{P}(\neg B) \Leftrightarrow \mathbb{P}(A)(1-\mathbb{P}(B)) > \mathbb{P}(A|\neg B)\mathbb{P}(\neg B),$$

where the first inequality follows from (3) and the last inequality can be divided by $\mathbb{P}(\neg B) = 1 - \mathbb{P}(B) > 0$ to obtain $P(A) > P(A|\neg B)$, which together with (3) implies $\mathbb{P}(A|B) > \mathbb{P}(A|\neg B)$. "$\Leftarrow$": Because $\mathbb{P}(A|\neg B) < \mathbb{P}(A|B)$, this follows from

$$\mathbb{P}(A) = \mathbb{P}(A|B)\mathbb{P}(B) + \mathbb{P}(A|\neg B)\mathbb{P}(\neg B) < \mathbb{P}(A|B)[\mathbb{P}(B) + \mathbb{P}(\neg B)] = \mathbb{P}(A|B).$$

□

However, if the ratio $\mathbb{P}(A|B)/\mathbb{P}(A|\neg B)$ were to be taken as a point of departure for evaluating cross-impacts, the resulting ratios would be asymmetric and consequently the number of parameters in the model would become much higher. Moreover, it could be cognitively more challenging for the respondent to specify statements involving comparisons in which the event A is conditioned on the *non-occurrence* of B.

The interpretation of cross-impacts in (4) implies that

$$\frac{C^{ij}_{kl}}{C^{ij}_{kl'}} = \frac{p^{ij}_{kl}}{p^i_k p^j_l} \times \frac{p^i_k p^j_{l'}}{p^{ij}_{kl'}} = \frac{p^{i|j}_{k|l}}{p^{i|j}_{k|l'}}. \tag{6}$$

Thus, the ratio between two different cross-impact terms provides information about "How many times more probable is the occurrence of x^i_k when x^j_l occurs, as opposed to when $x^j_{l'}$ occurs?" (cf. the discussion of Bayes factors; Kass & Raftery, 1995).

More generally, an important benefit of this probabilistic interpretation of cross-impact assessments is that the accuracy of such statements can be tested empirically, for instance by carrying out experiments with controlled subjects or by revisiting earlier cross-impact studies and examining how frequently the observed outcomes match those implied by the stated cross-impact ratios. These kinds of empirical studies help assess to what extent the statements may need to be calibrated to ensure a better fit with empirically observed marginal and conditional probabilities (see, e.g., Hora, 2007; O'Hagan et al., 2006).

2.2 | Relationship between cross-impact statements and scenario probabilities

The elicitation of statements about the ratio (4) for several pairs of uncertainty factors and their outcomes constitutes an approach to the elicitation of a dependency structure Werner et al. (2017). In this process, it is possible to employ discrete scales which translate numerical or verbal statements about how strongly the outcomes being assessed enforce each other into corresponding ranges of probability ratios (see Theil, 2002). To ensure the validity of assessments, these translations need to be properly justified and clearly communicated so that they can be understood by respondents.

The cross-impact ratio between the outcomes indexed by k, l of factors i and j is related to the joint probability distribution $p(\cdot): S \mapsto [0, 1]$ through

$$C^{ij}_{kl} = \frac{p^{ij}_{kl}}{p^i_k p^j_l} = \frac{\sum_{s \in S^{ij}_{kl}} p(s)}{\left(\sum_{s \in S^i_k} p(s)\right)\left(\sum_{s \in S^j_l} p(s)\right)}, \tag{7}$$

where $p(s) := \mathbb{P}(s), s \in S$ denote scenarios probabilities, the set S^{ij}_{kl} contains those scenarios in which the outcomes of factors i and j are x^i_k and x^j_l, respectively, and the set S^i_k consists of those scenarios in which the outcome of the ith uncertainty factor is x^i_k (and similarly for the scenario set S^j_l).

We assume that all outcomes of uncertainty factors occur with a probability that is strictly positive, that is, $p^i_k > 0, i = 1, \ldots, n$, $k = 1, \ldots, n_i$. This assumption is plausible, because otherwise the "impossible" outcome x^i_k such that $p^i_k = 0$ could be removed from the analysis. Technically, this assumption can be introduced through the constraint $p^i_k \geq \varepsilon$ where $\varepsilon > 0$ is a very small number.

Because the expression (7) is nonlinear in $p(s), s \in S$ with quadratic terms in the denominator, it is not possible to convert upper and lower bounds on this ratio into linear constraints on scenario probabilities. This is in contrast to bounds on marginal or conditional probabilities which both can be modeled through linear constraints on scenario probabilities (see Salo & Bunn, 1995).

The expression (7) can be written in matrix notation as follows. Let the set of all n scenarios be $S = \{s_1, s_2, \ldots, s_{|S|}\}$ and let $|S|$ denote the cardinality of S, that is the total number of scenarios. Furthermore, let the vector $p \in \mathbb{R}^{|S|}$ contain all the scenario probabilities so that probability of the ith scenario is $p_i = \mathbb{P}(s_i)$.

To link scenarios to the specific outcomes of uncertainty factors, we employ $m \times 1$ dimensional binary vectors $\sigma^i_k \in \{0, 1\}^{|S|}$ so that the mth element of this vector is 1 if the realization of the ith uncertainty factor in scenario s_m is x^i_k and zero otherwise. Then, the probability of the outcome x^i_k can be derived from the joint probability distribution over scenarios through

$$p^i_k = \mathbb{P}\left(X^i = x^i_k\right) = \sum_{j=1}^{|S|} p_j \left(\sigma^i_k\right)_j = \left(\sigma^i_k\right)^{\intercal} p, \tag{8}$$

where $^{\intercal}$ denotes the transpose of a matrix. The conditional probability $p^{i|j}_{k|l}$ in (4), in turn, can be written as

$$p^{i|j}_{k|l} = \frac{p^{ij}_{kl}}{p^j_l} = \frac{\left(\sigma^i_k \circ \sigma^j_l\right)^{\intercal} p}{\left(\sigma^j_l\right)^{\intercal} p}, \tag{9}$$

where the Hadamard product $\circ$ is defined as $(\sigma^i_k \circ \sigma^j_l)_m = (\sigma^i_k)_m (\sigma^j_l)_m$, $m = 1, \ldots, |S|$. Thus, the entry for the mth scenario in the vector $\sigma^i_k \circ \sigma^j_l$ is equal to 1 if and only if the outcomes of the ith and jth are equal to x^i_k and x^j_l. Placing lower and upper bounds $p^i_k \in \left[\underline{p}^i_k, \bar{p}^i_k\right]$ on the expression (8) leads to linear constraints on scenario probabilities. The linear fractional expression in (9) is the ratio between sums of those scenario probabilities which are picked by the vectors $\sigma^i_k \circ \sigma^j_l$ and σ^j_l, respectively. Thus, bounding this ratio through bounds $p^{i|j}_{k|l} = \left[\underline{p}^{i|j}_{k|l}, \bar{p}^{i|j}_{k|l}\right]$ can be transformed into linear constraints by multiplying these bounds by the denominator $(\sigma^j_l)^{\intercal} p$. For instance, the constraint $p^{i|j}_{k|l} \leq \bar{p}^{i|j}_{k|l}$ is equivalent to $(\sigma^i_k \circ \sigma^j_l)^{\intercal} p \leq \bar{p}^{i|j}_{k|l}[(\sigma^j_l)^{\intercal} p]$.

The cross-impact ratio (7) can be written as

$$C^{ij}_{kl} = \frac{p^{ij}_{kl}}{p^i_k p^j_l} = \frac{\left(\sigma^i_k \circ \sigma^j_l\right)^{\intercal} p}{\left[\left(\sigma^i_k\right)^{\intercal} p\right]\left[\left(\sigma^j_l\right)^{\intercal} p\right]}, \tag{10}$$

which is the same as the equality $C^{ij}_{kl}[(\sigma^i_k)^{\intercal} p][(\sigma^j_l)^{\intercal} p] = (\sigma^i_k \circ \sigma^j_l)^{\intercal} p$, which, in turn, is equivalent to the quadratic constraint

$$C^{ij}_{kl} \frac{1}{2} p^{\intercal} \mathbf{Q}^{ij}_{kl} p - \left(\sigma^i_k \circ \sigma^j_l\right)^{\intercal} p = 0, \tag{11}$$

where $\mathbf{Q}^{ij}_{kl} = (\sigma^i_k (\sigma^j_l)^{\intercal} + \sigma^j_l (\sigma^i_k)^{\intercal})$ is a symmetric matrix.

Thus, the modeling of cross-impact statements about the (4) leads to *quadratic* constraints on the scenario probabilities. As in the case of marginals and conditionals, these constraints can be introduced by eliciting lower and upper bounds on the cross-impact terms (i.e., $C^{ij}_{kl} \in \left[\underline{C}^{ij}_{kl}, \bar{C}^{ij}_{kl}\right]$, $\underline{C}^{ij}_{kl} \leq \bar{C}^{ij}_{kl}$) which impose inequality constraints on the underlying scenario probabilities. Yet, because the matrix $\mathbf{Q}^{ij}_{kl}$ can be indefinite, this set of scenario probabilities may be nonconvex, making it computationally more challenging to explore the implications of cross-impact statements for probabilistic inference. There are, however, specialized algorithms for optimization problems with quadratic terms in the objective function or in the constraints (see, e.g., Audet et al., 2000). These algorithms have been incorporated in commercial optimization solvers which are capable of handling problems such as the example in Section 4.

2.3 | Consistency implications of probabilistic statements

Because cross-impact statements refer to the same set of underlying scenario probabilities based on the ratio (4), these statements are interdependent in the sense that a given statement about any cross-impact term imposes constraints on the values of the cross-impact term for other pairs of uncertainty factors and their outcomes. One such example is the ratio (6) which connects pairs of cross-impact terms.

Specifically, if the implications of the earlier statements are not observed when introducing new ones, the constraints implied by the new statements may conflict with the constraints derived from the earlier ones. In this case, there are no feasible scenario probabilities which satisfy the full set of constraints that are associated with all the earlier and the newer statements.

To prevent this possibility, we strongly recommend that the consistency of the model should be maintained throughout the elicitation process so that new statements are introduced only on the condition that the resulting augmented set of constraints continues to be satisfied by at least some feasible scenario probabilities. One reason for this is that resolving a complex set of mutually inconsistent constraints can pose conceptual and computational difficulties. That is, it would call for the identification of those statements that are more "wrong" than others, leading to either the removal or relaxation of constraints that are associated with earlier statements.

In practice, the consistency of the statements can be supported so that the expression for the new statement to be added (i.e., marginal probability (8), conditional probability (9), or cross-impact statement (10)) is employed as the objective function which is then minimized and maximized subject to the constraints implied by the earlier statements. That is, the interval defined by these lower and upper consistency bounds indicates for which values the new statement is consistent with the earlier ones. The new statement will eliminate some previously feasible scenario probabilities from further consideration if and only if it excludes some values from the interval defined by the consistency bounds.

For example, consider the situation in which the cross-impact term C_{kl}^{ij} is about to be specified in terms of its lower and upper bounds $\left[\underline{C}_{kl}^{ij}, \overline{C}_{kl}^{ij}\right]$. Then, if the minimum of the difference on the left side in (11) is strictly positive for the cross-impact term $\underline{C}_{kl}^{ij}$, the new constraint will be excessively restrictive in that none of the feasible probabilities will satisfy the constraint based on $\underline{C}_{kl}^{ij}$. Conversely, if the maximum of this difference is strictly negative for the constraint based on $\overline{C}_{kl}^{ij}$, this upper bound is too restrictive. In this way, optimization problems can be solved to ensure the consistency of statements.

There are also further consistency checks that can be readily carried out by checking inequality expressions. First, note that the equality $\sum_{l=1}^{n_j} p_{k|l}^{i|j} p_l^j = p_k^i$ can be divided by p_k^i to obtain

$$\sum_{l=1}^{n_j} \frac{p_{k|l}^{i|j}}{p_k^i} p_l^j = \sum_{l=1}^{n_j} C_{kl}^{ij} p_l^j = 1,$$

which shows that the *probability-weighted average of cross-impact terms on any row of the cross-impact matrix for uncertainty factorsiandjmust equal one*. Thus, if $\underline{C}_{kl}^{ij}, \overline{C}_{kl}^{ij}$ are the lower and upper bounds on the next cross-impact ratio C_{kl}^{ij} which is being elicited, there must exist some feasible vector *p* of scenario probabilities such that the corresponding marginal probabilities p_l^j satisfy the inequalities $\sum_{l=1}^{n_j} \underline{C}_{kl}^{ij} p_l^j \leq 1 \leq \sum_{l=1}^{n_j} \overline{C}_{kl}^{ij} p_l^j$. Similarly, examining the marginals p_l^j leads to the equality $\sum_{k=1}^{n_i} C_{kl}^{ij} p_k^i = 1$ so that the probability-weighted average of cross-impact terms in any *column* must be equal to one. Thus $\sum_{k=1}^{n_i} \underline{C}_{kl}^{ij} p_k^i \leq 1 \leq \sum_{k=1}^{n_i} \overline{C}_{kl}^{ij} p_k^i$ for any $l = 1, \ldots, n_j$.

Even further relationships between the marginal and conditional probabilities and the cross-impact ratios can be established. For example, because $\max\{p_k^i, p_l^j\} \leq 1$, it follows that $\overline{C}_{kl}^{ij} \geq C_{kl}^{ij} = p_{kl}^{ij}/(p_k^i p_l^j) \geq p_{kl}^{ij}/\min\{p_k^i, p_l^j\}$ and hence $p_{kl}^{ij} \leq \overline{C}_{kl}^{ij} \min\{p_k^i, p_l^j\}$. Thus, if the upper bound on the cross-impact term is small, then the probability p_{kl}^{ij} of the joint event will be low relative to the marginal probabilities. In the same vein, using the inequality $p_{kl}^{ij} \leq \min\{p_k^i, p_l^j\}$ gives $\underline{C}_{kl}^{ij} \leq C_{kl}^{ij} = p_{kl}^{ij}/(p_k^i p_l^j) \leq \min\{p_k^i, p_l^j\}/(p_k^i p_l^j) = 1/\max\{p_k^i, p_l^j\}$ so that $\max\{p_k^i, p_l^j\} \leq 1/\underline{C}_{kl}^{ij}$. In other words, having a very large lower bound on the cross-impact term will place an upper bound on the marginal probabilities.

If consistency bounds are not systematically employed in the elicitation process, there are strategies which can be applied to preserve the consistency of the model. That is, if it is only the most recently elicited statement that is found to be inconsistent with the earlier statements, then it is possible to backtrack by omitting this statement from consideration. More constructively, the respondent can also be asked to revise the lower and upper bounds of this statement so that consistency is preserved. In principle, one could also seek to identify those subsets of statements that are mutually consistent and contain as many statements as possible (for a discussion of analogous approaches in the case of constraints on marginal and conditional probability statements, see Salo & Bunn, 1995). However, in the case of cross-impact statements, this strategy would call for a considerable amount of computational effort and, in addition, require that the respondent is prepared to indicate which one(s) of the earlier statements should be omitted.

We also note that the assessment of inconsistencies in nonprobabilistic cross-impact analysis differs from our approach. In the CIB method, for example, all cross-impact statements are elicited at the outset, whereafter an algorithm is applied to identify the scenarios that satisfy the consistency criterion. By construction, the application of this criterion in the CIB method presumes that all the statements have been elicited (i.e., it is not possible to exclude inconsistent scenarios based on a subset of cross-impact statements). Also, because this criteria lacks a formal theoretical foundation, it appears that nonprobabilistic approaches in which the consistency of scenarios is not treated as a dichotomous "yes-no" criterion but, rather, quantified by providing a more systemic measure of consistency, are more defensible. One such approach is developed by Seeve and Vilkkumaa (2021) who generate sets of plausible scenarios which are diverse, too, as measured by how different the scenarios are from each other.

In this context, it is worth noting that "comprehensiveness" has different connotations in nonprobabilistic and probabilistic approaches. In nonprobabilistic approaches, comprehensiveness refers to the extent to which the set of generated scenarios represents the entire range of possible futures (which, as a criterion, does not require that *all* the possible futures would have to be generated). In probabilistic approaches, and especially in the context of safety-critical systems, comprehensiveness commonly refers to the extent to which the residual uncertainties concerning the attainment of the safety requirements permit conclusive statements about the safety of the system (for a review and discussion, see Tosoni et al., 2018).

Furthermore, we remark that "consistency" has a somewhat different meaning in the CIB method than in our approach. In the former, consistencies are associated with individual entries of the cross-impact matrix (with 3 indicating *strong consistency* and −3 representing *strong inconsistency*) as well as with those *scenarios* that fulfill the criterion in Equation (1). In our approach, consistencies refer to *sets of statements* such that the corresponding constraints are fulfilled by some joint probability distribution over the set of all possible scenarios. That is, the scenarios are not treated as inconsistent as such but, rather, they are more or less probable, depending on the logical implications of the elicited probability statements. Also from this perspective of offering insights into what these statements signify, there are advantages to maintaining the consistency of the probability model, because this permits many kinds of probabilistic inferences, such as deriving bounds on those marginal and conditional probabilities that have not yet been elicited.

3 | CONDITIONING CONSEQUENCES ON SCENARIOS

In risk assessment, the aim is to characterize the magnitude of risks, as measured by the severity and probability of harmful consequences. These consequences can differ considerably in terms of what kinds of impacts they pertain to (e.g., human casualties, environmental damages, financial losses).

We first consider the situation where these consequences are represented by a real-valued random variable Z (e.g., amount of released radioactivity from a nuclear facility) whose realization depends on which one of the $|S|$ scenarios occurs. Because the scenarios are mutually exclusive and collectively exhaustive, the probability for the event that the consequences exceed a given threshold level $\theta \in \mathbb{R}$ (e.g., a regulatory limit) is obtained by conditioning Z on these scenarios $s \in S$ so that

$$\mathbb{P}(Z > \theta) = \sum_{s \in S} \mathbb{P}(Z > \theta | s)\mathbb{P}(s). \tag{12}$$

In risk assessment, one relevant rationale for the development of scenarios is that the approach of assessing the conditional probabilities $\mathbb{P}(Z > \theta | s)$ for the different scenarios separately can lead to a more structured and defensible elicitation process than seeking to obtain a single holistic estimate $\mathbb{P}(Z > \theta)$ (for an overview of structured elicitation methods, see Dias et al., 2018).

The expression (12) can be also employed to shed light on the question about which one(s) out of further candidates for additional uncertainty factors $X^{n+1}, X^{n+2}, \ldots$ should be introduced to complement the n uncertainty factors $X^1, \ldots, X^n$, on the basis of which scenarios have already been formulated. Toward this end, the scenario-based conditioning of $\mathbb{P}(Z > \theta | s)$ can be extended to include the additional uncertainty factor X^{n+1} so that

$$\mathbb{P}(Z > \theta | s) = \sum_{x_k^{n+1} \in S^{n+1}} \mathbb{P}\Big(Z > \theta | s, X^{n+1} = x_k^{n+1}\Big)\mathbb{P}\Big(s \wedge \Big\{X^{n+1} = x_k^{n+1}\Big\}\Big).$$

In particular, this expression suggests that the inclusion of the additional uncertainty factor X^{n+1} is unlikely to be very useful if (i) the conditional probabilities $\mathbb{P}(Z > \theta | s, X^{n+1} = x_k^{n+1})$ are the same for different outcomes $x_k^{n+1} \in S^{n+1}$ (i.e., the first term in the sum is the same for all outcomes of the uncertainty factor X^{n+1}) or if (ii) the factor X^{n+1} is perfectly correlated with any one of the n factors that are included in the scenarios $s \in S$ (i.e., there exists some other factor $X^i, i = 1, \ldots, n$ such that the outcomes of X^{n+1} are implied by the states of X^i). These conditions, together with an assessment of how much extra effort is required to elicit the additional parameters $\mathbb{P}(Z > \theta | s, X^{n+1} = x_k^{n+1})$ and $\mathbb{P}(s \wedge \{X^{n+1} = x_k^{n+1}\})$, help evaluate which additional uncertainty factors should be included in the analysis.

Furthermore, the expression (12) implies that if some scenarios are omitted from the sum on the right side, the assessed probability for the event $Z > \theta$ will be lower than the actual probability, unless this omission is compensated through an upward adjustment in the other terms in the sum. Furthermore, if the aim is to establish a conservative upper bound for $\mathbb{P}(Z > \theta)$, then the estimates employed for the terms $\mathbb{P}(Z > \theta | s)$ should be upper bounds on these scenario-specific probabilities.

The expression (12) can also be generalized to situations where Z is not necessarily real-valued but takes on values in the set of possible consequences $\mathbb{C}$. An appropriate disutility function $U : \mathbb{C} \mapsto \mathbb{R}$ can then be defined so that the value of this function is highest for the least preferred consequences and lowest for most preferred consequences. Such a disutility function can be also defined to characterize the probability with which these consequences will be unacceptable. That is, let the set $\mathbb{C}^{\text{fail}}$ consist of all unacceptable consequences and define the disutility function so that

$$U(Z) = \begin{cases} 1, & Z \in \mathbb{C}^{\text{fail}} \\ 0, & Z \notin \mathbb{C}^{\text{fail}}. \end{cases}$$

For this disutility function, the expression $\sum_{s \in S} \mathbb{E}[U(Z)|s]p(s)$ gives the probability with which the consequences will be unacceptable. More generally, we assume that the risk assessment process is required to provide conservative estimates for the expressions

$$\mathbb{E}[Z] = \sum_{s \in S} \mathbb{E}[Z|s]p(s), \tag{13}$$

$$\mathbb{E}[U(Z)] = \sum_{s \in S} \mathbb{E}[U(Z)|s]p(s), \tag{14}$$

where in (13) the term Z representing consequence is assumed to be real-valued and the disutility function in (14) makes it possible to handle other types of consequences as well.

Using the notations $u(s) = \mathbb{E}[U(Z)|s]$, the above formulations can be combined with the results of the preceding section to state the following optimization problem

$$\begin{aligned} \max/\min_{p(s)} \quad & \sum_{s \in S} u(s)p(s) \\ \text{subject to} \quad & \sum_{s \in S} p(s) = 1, \\ & p \geq 0, \end{aligned} \tag{15}$$

plus all the constraints that correspond to the elicited statements about the marginal probabilities, conditional probabilities, and cross-impact terms. Thus, lower and upper bounds for the risk level can be estimated by solving the optimization problem as a minimization and a maximization, respectively, of the objective function.

Building on the above, the main phases of probabilistic cross-impact analysis for assessing risks can now be outlined as follows:

1. Define the scenarios $s \in S$ by specifying the uncertainty factors and their possible outcomes.
2. Assess bounds for the expected scenario-specific consequences $\mathbb{E}[Z|s]$ or their expected disutilities $\mathbb{E}[U(Z)|s]$.
3. Obtain information about the joint probability distribution over scenarios by eliciting cross-impact statements about the ratio (4) and/or statements about the marginal and conditional probability distributions (see Salo & Bunn, 1995). These statements can be elicited by employing interval valued statements defined by lower and upper bounds.
4. Compute lower and upper bounds for the aggregate risk level (as expressed in (12), (13), or (14)) based on information about corresponding scenario-specific expectations and the joint probability distribution over scenarios.
5. Once the maximum tolerable risk level has been determined, assess the risk management implications of the available information by considering the following possibilities (Tosoni et al., 2019):
   - If the upper bound of the aggregate risk level, obtained by maximizing (15), is below the maximum tolerable risk level, the system can be deemed safe.
   - If the lower bound of the aggregate risk level, obtained by minimizing (15), is higher than the maximum tolerable risk level, the system can be deemed unsafe.
   - Otherwise, return to steps 2 and 3 to obtain additional information with the aim of deriving tighter bounds on the aggregate risk level.

From the viewpoint of data analysis and generation, solving the problem (15) presumes that the expected scenario-specific disutilities $u(s)$, $s \in S$ are available for all scenarios. There are, however, problem contexts in which estimates about these disutilities can be generated with the help of computational models, as illustrated by the example in the next section. The maximization problem (15) can also be solved based on conservative upper bound estimates about these disutilities. One can also explore just how large these disutilities would have to be so that the maximum tolerable risk level would be reached.

Because the cross-impact statements are interpreted as constraints on the joint probabilities, it is conceptually and computationally straightforward to integrate the use of such statement in Monte Carlo simulations in which vectors representing joint probabilities are generated. That is, computational results reflecting cross-impact statements can be produced by retaining only those probability vectors that satisfy the constraints implied the cross-impact statements. In particular, this makes it possible to benefit from cross-impact statements when using other approaches for the exploration dependencies in safety risk models (see, e.g., Harrison & Cheng, 2011).

4 | CASE STUDY

The risk assessment of nuclear waste management facilities is an important application context of scenario analysis (Tosoni et al., 2018). In this context, the uncertainty factors consist of so-called FEPs which include, for instance, physical and chemical variables that affect the lifetime of the facility and its surrounding environment. The FEP outcomes can be represented through discretized states such as *low*, *medium*, and *high*.

In this section, we revisit the case study (Tosoni et al., 2019) on the nuclear waste repository at Dessel (Belgium) in which the Bayesian network in Figure 3 was developed to represent dependencies between nine FEPs. As shown in Table 1, there are two possible outcomes for the first five FEPs while the two last ones have three possible outcomes.

In this setting, scenarios are defined as combinations of outcomes for each FEP. Thus, for example, there is a scenario which represents the following combination of FEP states: a *beyond-design-basis* Earthquake (BDBE), *low* Water flux, *micro* crack Aperture, *low* Diffusion coefficient, *low* Distribution coefficient, *slow* Chemical degradation, *fast* Concrete degradation, *slow* Monolith degradation, and *low* Hydraulic conductivity. Given the nine FEPs and their two or three outcomes, the total number of scenarios is $2^7 \times 3^2 = 1152$.

The scenarios differ from each other in terms of how probable it is that radioactive particles will be released into the environment, causing human exposure to radiation. For each scenario, this impact is quantified by the conditional probability that the subsequent dose rate to humans exceeds a predefined safety threshold level. Aggregating these conditional probabilities over all scenarios based on (12) thus gives an estimate about the radiological risk, which is measured by the total probability with which this threshold θ is violated.

For each scenario s, the corresponding conditional probability $\mathbb{P}(Z > \theta|s)$ in (12) of violating the threshold θ was computed as the average of three numbers, that is, (i) the prior value in Tosoni et al. (2020) and (ii) the lower and upper bounds in Tosoni et al. (2019). This approach was adopted, because it serves to illustrate how results concerning the total violation probability $\mathbb{P}(Z > \theta)$ in (12) changes as a result of providing additional information about the probabilities. These conditional violation probabilities are not reported here due to the large number of scenarios, but they are available from the authors upon request. For instance, the conditional violation probability for the scenario described in the second paragraph of this section was 0.678.

In the following illustrative analysis, we build on the model and data in papers Tosoni et al. (2019, 2020) which represent the nuclear waste repository as a Bayesian network (Pearl & Russel, 2003). In this network, the nodes represent the FEPs, whereas directed arcs indicate cause dependencies between the FEPs. The uncertainties associated with the FEP outcomes are modeled as the feasible sets of marginal and conditional probabilities (Tosoni et al., 2019).

Specifically, we consider three steps in which increasingly detailed information about scenario probabilities are provided. The first step uses only marginal probabilities of FEP outcomes. In the second step, the dependencies between those FEPs which are linked by arcs in the Bayesian network are approximated with cross-impact statements. In the third step, it is stated that the six FEPs in Figure 3 (i.e., Water flux, Earthquake, Crack aperture, Diffusion coefficient, Distribution coefficient, Chemical degradation) from which there are only outgoing arcs are almost independent. This statement is introduced by allowing the cross-impact ratio (4) to assume value in the

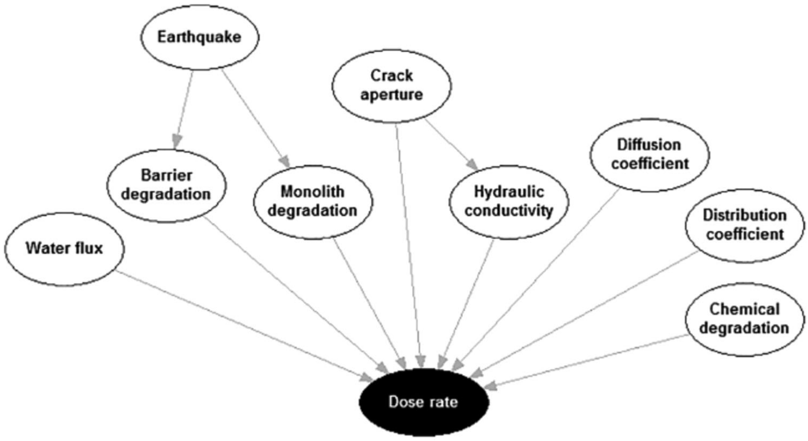

FIGURE 3 The Bayesian network for the case study (Tosoni et al., 2019)

interval [0.9950–1.0050]. Note that this assumption is weaker than the full independence assumption which is embedded in the structure of the Bayesian network and corresponds to the requirement that all cross-impacts between the outcomes of these six FEPs are equal to 1. Thus, the introduction of the relatively narrow interval [0.9950–1.0050] helps explore how the results would change if it were to be the case that the Bayesian network in Figure 3 is *not* a valid model of the dependencies between the FEPs. Furthermore, because the statements in second step do not yet limit these dependencies, the introduction of these intervals in the third step provides a significant amount of additional information. This leads to much tighter constraints on the scenario probabilities so that a reduction in the violation probability $\mathbb{P}(Z > \theta)$ can be expected.

For the first step, the lower and upper bounds for the marginal probabilities of FEP outcomes in Table 1 were computed by sampling the feasible sets of marginal and conditional probabilities in the Bayesian network, leading to corresponding sample distributions over FEP outcomes. The marginals in Table 1 were taken from these distributions by employing their 5% and 95% quantiles.

For the second step, the characterization of dependencies between selected pairs of FEP outcomes was also based on the model in Tosoni et al. (2019) as above, except that the sample distributions were established for the cross-impact ratio in (7) (rather than for the marginal probability distributions). Moreover, the bounds for cross-impact terms were established by using the more conservative 0.5% and 99.5% quantiles (as opposed to 5% and 95% quantiles) to allow for more imprecision in the characterization of cross-impacts. The resulting bounds on the cross-impact ratios are reported in Tables 2 and 3.

Looking at the ratios in Table 2, it is instructive to see that if a *major* outcome of the FEP Earthquake does occur, the probability of *fast* Barrier degradation becomes much higher (i.e., [1.5755–6.9785] times) in comparison with the situation where there is no information about the probability of an Earthquake. On the other hand, if the outcome for the Earthquake is *BDBE* (i.e., of a lower magnitude than a major earthquake, but still beyond what the repository barriers are designed to withstand), the probability of *slow* Barrier degradation will grow, albeit marginally. This is in keeping with the recognition that a *major* Earthquake can have an impact on the speed of Barrier degradation; but its absence does not have a comparable impact.

TABLE 1 FEPs and their outcomes in Tosoni et al. (2019) and corresponding on bounds marginal probabilities

FEP	Outcome	Probability bounds
Earthquake	BDBE	[0.9912–0.9950]
	Major	[0.0050–0.0088]
Water flux	Low	[0.6525–0.8428]
	High	[0.1572–0.3475]
Crack aperture	Micro	[0.8148–0.8874]
	Macro	[0.1126–0.1852]
Diffusion coefficient	Low	[0.5209–0.7275]
	High	[0.2725–0.4791]
Distribution coefficient	Low	[0.5215–0.7268]
	High	[0.2732–0.4785]
Chemical degradation	Fast	[0.5361–0.6694]
	Slow	[0.3306–0.4639]
Barrier degradation	Fast	[0.0787–0.2337]
	Slow	[0.7663–0.9213]
Monolith degradation	Very fast	[0.0293–0.2678]
	Fast	[0.0594–0.2695]
	Slow	[0.4627–0.9114]
Hydraulic conductivity	Low	[0.5993–0.7066]
	Medium	[0.2016–0.2715]
	High	[0.0872–0.1342]

Abbreviation: BDBE, beyond-design-basis earthquake; FEPs, features, events, and processes.

TABLE 2 Bounds on the cross-impact ratios for pairs of outcomes for the FEPs Earthquake, Barrier degradation, and Monolith degradation

		Barrier degradation		Monolith degradation		
		Fast	**Slow**	**Very fast**	**Fast**	**Slow**
Earthquake	BDBE	[0.9544–0.09963]	[1.0011–1.0036]	[0.9950–1.0050]	[0.9329–0.9982]	[0.9975–1.0052]
	Major	[1.5755–6.9785]	[0.5660–0.8174]	[0.9950–1.0050]	[1.2769–10.1853]	[0.3406–1.3606]

Abbreviations: BDBE, beyond-design-basis earthquake; FEPs, features, events, and processes.

TABLE 3 Bounds to the cross-impact ratios for pairs outcomes of the FEPs Crack aperture and Hydraulic conductivity

		Hydraulic conductivity		
		Low	**Medium**	**High**
Crack	Micro	[1.0896–1.1880]	[0.4941–0.7490]	[0.9950–1.0050]
Aperture	Macro	[0.1628–0.3017]	[2.5666–3.9985]	[0.9950–1.0050]

Abbreviation: FEPs, features, events, and processes.

The bounds in Tables 2 and 3 specify no restrictions on dependencies between the six FEPs which are independent in Figure 3 as they have only outgoing arcs in this Bayesian network. Thus, the independence between these six FEPs is introduced in the third step. As noted above, however, this independence assumption is quite strong, so that we relax it by allowing for minor deviations from independence by bounding the cross-impact ratios to the interval [0.995–1.005]. Moreover, in Tables 2 and 3 there are two columns (i.e., *very fast* Monolith degradation in Table 2, *high* Hydraulic conductivity in Table 3) in which the independence assumption contained in the Bayesian data has been relaxed similarly.

Based on the probability information for the three steps above, the following conservative upper bounds for the level of radiological risk can now be computed by solving the maximization problem (15) subject to the corresponding constraints on scenario probabilities.

1. *Marginals only*: When there is information about the marginals only, the upper bound on the maximum level of risk is 0.576.
2. *Cross-impacts bounds for arcs between FEPs in* Tables 2 and 3: When the constraints based on these bounds are added to the information in the first step, the upper bound is reduced to 0.571.
3. *Cross-impact bounds for independent FEPs*: When the narrow intervals [0.995–1.005] are introduced for pairs of outcomes for independent FEPs in the Bayesian network, the upper bound becomes 0.427.

The results are summarized in Table 4. The greatest reduction in the upper bound is attained as a result of introducing the assumption of near-independence when moving from the second step to the third. This can be explained by noting that the number of such constraints is high (i.e., lower and upper bound constraints for every combination of outcomes for all pairs of the six FEPs) and because these intervals are relatively tight. This can be contrasted with the shift from the first step to the second step which leads to a much smaller reduction in the total violation probability.

TABLE 4 Upper bounds on the risk level for different settings of probabilistic information

	Setting	1	2	3
Constraints	1. Marginals	✓	✓	✓
	2. CI ratios for designated FEP dependencies (Tables 2, 3)	✗	✓	✓
	3. CI ratios for independent FEPs	✗	✗	✓
	Upper bound on risk level	0.576	0.571	0.427

Abbreviation: FEPs, features, events, and processes.

More generally, this example shows how probabilistic cross-impact analysis can be interfaced with other models. Specifically, scenario-specific estimates concerning radiological risk were inferred from Tosoni et al. (2019, 2020). Parameters of the Bayesian network (Tosoni et al., 2019) were employed to generate information about the marginal probabilities. Analogously, information about conditional dependencies was provided through cross-impact ratios stated in terms of lower and upper bounds. We emphasize that all this information about probabilities and dependencies could have been introduced directly without explicit reference to the Bayesian network (which has been employed as a useful tool for generating such information). This notwithstanding, we stress that the numerical results are illustrative and do not provide any indications as to the safety of the nuclear waste repository at Dessel.

5 | DISCUSSION AND CONCLUSIONS

In this paper, we have considered the limitations of nonprobabilistic cross-impact analyses in risk management and, specifically, in the risk assessment of safety critical systems for which the aim is produce conservative estimates that provide an upper bound on the overall risk level. Importantly, we have shown that instead of limiting attention to the most consistent scenarios only, it is pertinent to account for all the scenarios that can make a nonnegligible contribution to the overall risk level, even if some of these scenarios are quite improbable. That is, neglecting these scenarios may lead to risk estimates which are too small, as the actual risk will be higher than what is suggested by the analysis. This, in turn, may lead to the selection of inadequate and insufficient risk mitigation actions.

We have also advocated the probabilistic interpretation of cross-impacts, because this helps establish precise and empirically testable mappings between the qualitative verbal expressions employed in the elicitation process and their numerical counterparts. This interpretation also makes it possible to integrate the scenario process with other approaches for analyzing probabilistic inputs, for instance by carrying out statistical analyses or by synthesizing them with judgmental forecasts (see, e.g., G. Wright et al., 2009). Furthermore, probabilistic models are appealing not least because they can be adapted to assess the attractiveness and effectiveness of insurance as one of the quantitative risk management options.

We have also developed a probabilistic cross-impact method which is capable of accommodating and synthesizing many kinds of probability elicitation statements (including both marginal and conditional probabilities as well as cross-impacts statements). All these statements are converted into corresponding linear or quadratic constraints in the optimization models which can be solved to (i) guide the elicitation of further statements which are consistent with the statements that have been elicited earlier and (ii) compute lower and upper bounds on the overall risk level at any stage of the elicitation process. Results such as these are useful for reaching conclusions about the safety of the system, which provides support for risk management decisions. There are also promising avenues for future work, for example by employing cross-impact statements together with other methods for assessing dependencies and their impacts (see, e.g., Harrison & Cheng, 2011). One could also assess how the cross-impact statements and therefore scenario probabilities, too, would be impacted by alternative risk management actions. This would make it possible to accommodate endogenously dependent scenario probabilities (for a case study with decision-dependent scenario probabilities, see Vilkkumaa et al., 2018).

ACKNOWLEDGMENTS

The constructive comments by Anna Leinonen, Jarmo Lehikoinen, Gotcheva Nadezhda, Raija Koivisto, Kari Rasilainen, and three referees are gratefully acknowledged. This study has been supported by the project "Systematic scenario methods for the assessment of overall safety" (SYSMET), funded by the Finnish Research Programme on Nuclear Waste Management (KYT2022) 2019–2022.

CONFLICT OF INTERESTS

The authors declare that there are no conflict of interests.

DATA AVAILABILITY STATEMENT

Data are available on request from the authors.

ORCID

Ahti Salo http://orcid.org/0000-0003-1526-7923

REFERENCES

Audet, C., Hansen, P., Jaumard, B., & Savard, G. (2000). A branch and cut algorithm for nonconvex quadratically constrained quadratic programming. *Mathematical Programming, Series A*, *87*, 131–152.

Aven, T., & Zio, E. (2011). Some considerations on the treatment of uncertainties in risk assessment for practical decisionmaking. *Reliability Engineering and System Safety*, *96*, 64–74.

Bedford, T., & Cooke, R. M. (2001). *Probabilistic risk analysis: Foundations and methods*. Cambridge University Press.

Bowman, G. (2016). The practice of scenario planning: An analysis of inter-and-intra-organizational strategizing. *British Journal of Management*, *27*(1), 77–96.

Bradfield, R., Wright, G., Burt, G., Cairns, G., & Heijden, K. V. D. (2005). The origins and evolution of scenario techniques in long range business planning. *Futures*, *37*(8), 795–812.

Bunn, D. W., & Salo, A. (1993). Forecasting with scenarios. *European Journal of Operational Research*, *68*(3), 291–303.

Carlsen, H., Eriksson, E. A., Dreborg, K. H., Johansson, B., & Bodin, O. (2016). Systematic exploration of scenario spaces. *Foresight*, *18*(1), 59–75.

Dias, L. C., Morton, A., & Quigley, J. (Eds.). (2018). *Elicitation: The science and art of structuring judgement* (Vol. 261). Springer International Series in Operations Research & Management Science.

Duperrin, J. C., & Godet, M. (1975). SMIC 74—A method for constructing and ranking scenarios. *Futures*, *7*(4), 302–312.

Fenton, N. E., Neil, M., & Caballero, J. C. (2016). Using ranked nodes to model qualitative judgments in Bayesian networks. *IEEE Transactions on Knowledge and Data Engineering*, *19*(10), 1420–1432.

Godet, M. (2000). How to be rigorous with scenario planning? *Foresight*, *2*(1), 5–9.

Goodwin, P., & Wright, G. (2010). The limits of forecasting methods in anticipating rare events. *Technological Forecasting and Social Change*, *77*(3), 355–368.

Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, *37*, 424–438.

Harrison, C., & Cheng, D. (2011). Modelling uncertainty in the RSSB safety risk model. *Safety and Reliability*, *31*, 19–33.

Helton, J. C., & Sallaberry, C. J. (2009). Conceptual basis for the definition and calculation of expected dose in performance assessments for the proposed high-level radioactive waste repository at Yucca Mountain, Nevada. *Reliability Engineering and System Safety*, *94*, 677–698.

Hirabayashi, Y., Mahendran, R., Koirala, S., Konoshima, L., Yamazaki, D., Watanabe, S., Kim, H., & Kanae, S. (2013). Global flood risk under climate change. *Nature Climate Change*, *3*, 816–821.

Hora, S. (2007). Eliciting probabilities from experts. In W. Edwards, R. Miles, Jr., & D. Von Winterfeldt (Eds.), *Advances in decision analysis: From foundations to applications* (pp. 129–153). Cambridge University Press.

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*(430), 773–795.

Kemp-Benedict, E., Carlsen, H., & Kartha, S. (2010). Large-scale scenarios as 'boundary conditions': A cross-impact balance simulated annealing (CIBSA) approach. *Technological Forecasting and Social Change*, *143*, 55–63.

Laitila, P., & Virtanen, K. (2016). Improving construction of conditional probability tables for ranked nodes in Bayesian networks. *IEEE Transactions on Knowledge and Data Engineering*, *28*(7), 1691–1706.

Liesiö, J., & Salo, A. (2012). Scenario-based portfolio selection of investment projects with incomplete probability and utility information. *European Journal of Operational Research*, *217*(1), 162–172.

Lord, S., Helfgott, A., & Vervoort, J. M. (2016). Choosing diverse sets of plausible scenarios in multidimensional exploratory futures technique. *Futures*, *77*(1), 11–27.

National Grid. (2020). *Network options assessment methodology review*. Retrieved March 23, 2021, from https://www.nationalgrideso.com/document/90851/download

O'Hagan, A., Buck, C. E., Daneshkhah, A., Eiser, J. R., Garthwaite, P. H., Jenkinson, D. J., Oakley, J. E., & Rakow, T. (2006). *Uncertain judgements: Eliciting experts' probabilities*. John Wiley & Sons.

Panula-Ontto, J., Luukkanen, J., Kaivo-oja, J., O'Mahony, T., Vehmas, J., Valkealahti, S., Björkqvist, T., Korpela, T., Järventausta, P., Majanne, Y., Kojo, M., Aalto, P., Harsia, P., Kallioharju, K., Holttinen, H., & Repo, S. (2018). Cross-impact analysis of Finnish electricity system with increased renewables: Long-run energy policy challenges in balancing supply and consumption. *Energy Policy*, *118*, 504–513.

Pearl, J. (2013). *Causality: Models, reasoning, and inference* (2nd ed.). Cambridge University Press.

Pearl, J., & Russel, S. (2003). Bayesian networks. In M. A. Arbib (Ed.), *Handbook of brain theory and neural networks* (2nd ed., pp. 157–160). MIT Press.

Pöyhönen, M. A., Hämäläinen, R. P., & Salo, A. (1997). An experiment on the numerical modeling of verbal ratio statements. *Journal of Multi-Criteria Decision Analysis*, *6*(1), 1–10.

Rausand, M. (2013). *Risk assessment: Theory, methods, and applications*. John Wiley & Sons.

Salo, A., & Bunn, D. W. (1995). Decomposition in the assessment of judgemental probability forecasts. *Technological Forecasting & Social Change*, *49*(1), 13–25.

Salo, A., Tosoni, E., & Zio, E. (2021). *Risk importance measures for scenarios in probabilistic risk assessment with Bayesian networks*. Manuscript submitted for publication.

Schoemaker, P. J. (1993). Multiple scenario development: Its conceptual and behavioural foundation. *Strategic Management Journal*, *14*(3), 193–213.

Scholz, R. W., & Tietje, O. (2001). *Embedded case study methods: Integrating quantitative and qualitative knowledge* (1st ed.). Sage Publications.

Schweizer, V. J. (2020). Reflections on cross-impact balances: A systematic method constructing global socio-technical scenarios for climate change research. *Climatic Change*, *162*, 1705–1722.

Science for Environment Policy. (2017). *The precautionary principle: Decision making under uncertainty*. Future Brief 18. Produced for the European Commission DG Environment by the Science Communication Unit, UWE, Bristol. http://ec.europa.eu/science-environment-policy

Seeve, T., & Vilkkumaa, E. (2021). Identifying and visualizing a diverse set of plausible scenarios for strategic planning. *European Journal of Operational Research* (forthcoming). https://doi.org/10.1016/j.ejor.2021.07.004

Stern, N. (2007). *The economics of climate change: The stern review*. Cambridge University Press.

Suppes, P. (1970). *A probabilistic theory of causality*. North-Holland Publishing Company.

Theil, M. (2002). The role of translations of verbal into numerical probability expressions in risk management: A meta-analysis. *Journal of Risk*, *5*(2), 177–186.

Tietje, O. (2005). Identification of a small reliable and efficient set of consistent scenarios. *European Journal of Operational Research*, *162*(1), 418–432.

Tosoni, E. (2021, August). *Novel methods of scenario analysis for the probabilistic risk assessment of nuclear waste storage and disposal facilities* (104/2021) [Doctoral Dissertations, Aalto University, Helsinki].

Tosoni, E., Salo, A., Govaerts, J., & Zio, E. (2019). Comprehensiveness of scenarios in the safety assessment of nuclear waste repositories. *Reliability Engineering and System Safety*, *188*, 561–573.

Tosoni, E., Salo, A., Govaerts, J., & Zio, E. (2020). Definition of the data for comprehensiveness in scenario analysis of near-surface nuclear waste repositories. *Data in Brief*, *31*, 105780.

Tosoni, E., Salo, A., & Zio, E. (2018). Scenario analysis for the safety assessment of nuclear waste repositories: A critical review. *Risk Analysis*, *38*(4), 775–776.

USEPA. (2014). *Risk assessment forum white paper: Probabilistic risk assessment methods and case studies* (EPA/100/R-14/004). Office of the Science Advisor, U.S. Environmental Protection Agency. https://www.epa.gov/sites/default/files/2014-12/documents/raf-pra-white-paper-final.pdf

USNRC. (2016). *Probabilistic risk assessment and regulatory decision making: Frequently asked questions* (NUREG-2201). Office of Nuclear Regulatory Research, U.S. Nuclear Regulatory Commission. https://www.nrc.gov/reading-rm/doc-collections/nuregs/staff/sr2201/index.html

Van Dorp, R. (2005). Statistical dependence through common risk factors: With applications in uncertainty analysis. *European Journal of Operational Research*, *161*, 240–255.

Vilkkumaa, E., Liesiö, J., Salo, A., & Ilmola-Sheppard, L. (2018). Scenario-based portfolio model for building robust and proactive strategies. *European Journal of Operational Research*, *266*(1), 205–220.

Weimer-Jehle, W. (2006). Cross-impact balances: A system-theoretical approach to cross-impact analysis. *Technological Forecasting & Social Change*, *73*, 334–361.

Weimer-Jehle, W. (2008). Cross-impact balances: Applying pair interaction systems and multi-value Kauffman nets to multidisciplinary systems analysis. *Physica A*, *387*, 3689–3700.

Weimer-Jehle, W., Vögele, S., Hauser, W., Kosow, H., Poganietz, W.-T., & Prehofer, S. (2020). Socio-technical energy scenarios: State-of-the-art and CIB-based approaches. *Climatic Change*, *162*, 1723–1741.

Werner C., Bedford, T., Cooke, R. M., Hanea, A. M., & Morales-Nápoles, O. (2017). Expert judgement for dependence in probabilistic modeling: A systematic literature review and future research directions. *European Journal of Operational Research*, *258*(3), 801–819.

Wiebe, K. Zurek, M., Lord, S., Brzezina, N., Gabrielyan, G., Libertini, J., Loch, A., Thapa-Parajuli, R., Vervoort, J., & Westhoek, H. (2018). Scenario development and foresight analysis: Exploring options to inform choices. *Annual Review of Environment and Resources*, *43*, 545–570.

Williamson, J. (2009). Probabilistic theories [of causality]. In H. Beebee, C. Hitchcock, & P. Menzies (Eds.), *The Oxford handbook of causation* (pp. 185–212). Oxford University Press.

Wright, A. (2005). The role of scenario as prospective sensemaking devices. *Management Decision*, *43*(1), 86–101.

Wright, G., Bolger, F., & Rowe, G. (2009). Expert judgement of probability and risk. In T. M. Williams, K. Samset, & K. J. Sunnevaag (Eds.), *Making essential choices with scant information* (pp. 213–229). Palgrave MacMillan.

Wright, G., Bradfield, R., & Cairns, G. (2013). Does the intuitive logics method-and its recent enhancements-produce "effective" scenarios. *Technological Forecasting and Social Change*, *80*(1), 631–642.

Wright, G., Saunders, C., & Ayton, P. (1988). The consistency, coherence and calibration of holistic, decomposed and recomposed judgemental probability forecasts. *Journal of Forecasting*, *7*, 185–199.

How to cite this article: Salo, A., Tosoni, E., Roponen, J., & Bunn, D. W. (2022). Using cross-impact analysis for probabilistic risk assessment. *Futures & Foresight Science*, 4, e103. https://doi.org/10.1002/ffo2.103

Publication IV

Juho Roponen and Ahti Salo. A probabilistic cross-impact methodology for explorative scenario analysis. Submitted to *Futures & Foresight Science*, December 2022.

ORIGINAL ARTICLE

A Probabilistic Cross-Impact Methodology for Explorative Scenario Analysis

Juho Roponen[1] | **Ahti Salo[1]**

[1]Department of Mathematics and Systems Analysis, Aalto-university, Espoo, Finland

Correspondence
Juho Roponen, Department of Mathematics and Systems Analysis, Aalto-university, 00076 Aalto, Finland
Email: firstname.lastname@aalto.fi

Funding information
MATINE - The Scientific Advisory Board for Defence of the Finnish Ministry of Defence.

As one of the approaches to scenario analysis, cross-impact methods provide a structured approach to building scenarios as combinations of outcomes for selected uncertainty factors. Although they vary in their details, cross-impact methods are similar in that synthesize expert judgments about probabilistic or causal dependencies between pairs of uncertainty factors and seek to focus attention on scenarios that can be deemed consistent. Still, most cross-impact methods do not associate probabilities with scenarios, which limits the possibilities of harnessing them in risk and decision analysis. Motivated by this recognition, we develop a cross-impact method that derives a joint probability distribution over all possible scenarios from probabilistically interpreted cross-impact statements. More specifically, our method (i) admits a broad range of probabilistic statements about the realizations of uncertainty factors, (ii) supports the process of eliciting such statements, (iii) synthesizes these judgments by solving a series of optimization models from which the corresponding scenario probabilities are derived. The resulting scenario probabilities can be used to construct Bayesian networks, which expands the range of analyses that can be carried out. We illustrate our method with a real case study on the impacts of 3D-printing on the Finnish Defense Forces. The scenarios, their probabilities and the associated Bayesian network resulting from this case study helped explore alternative futures and gave insights into how the Defence Forces could benefit from 3D-printing.

KEYWORDS
Scenario analysis, *Cross-impact analysis*, Probability estimation

1 | INTRODUCTION

Over the past decades, scenario analysis has established itself as one of the most widely employed approaches to support long-term planning and strategic management (Scholz and Tietje, 2002; Chermack, 2022). The need to prepare for alternative futures has long been recognized in planning contexts that involve major investments that call for long-term commitments and give rise to far-reaching consequences. The energy sector, car manufacturing, and national defense, among others, are examples of such contexts (Bunn and Salo, 1993; Ritchey, 2009). In recent years, growing concerns over climate change have continued to motivate global efforts on the development of climate scenarios (Duinker and Greig, 2007; O'Neill et al., 2020). Recently, the need for informative scenarios has been accentuated by the severe health and geopolitical crises that have affected companies, public organizations, and entire countries.

The broad range of scenario methods reflects differences in the many contexts in which scenarios are built. In particular, the characteristics of scenario methods span the range from purely qualitative (Bowman, 2016; Schwartz, 2012) to purely quantitative (Pereira et al., 2010; Siljander and Ekholm, 2018), with a rich array of methods that combine aspects of both (Godet, 1986; Kemp-Benedict, 2004; Kosow and Gaßner, 2008). In general, this diversity reflects differences in the requirements of different application contexts. In effect, the properties of the appropriate method depending on questions such as 1) What is the context of strategic analysis, and what are the issues at stake? 2) For what purposes are the scenarios created and to whom are they presented? 3) What are the time horizons involved? 4) What data (qualitative or quantitative) can be obtained to support the formulation of scenarios? Considering questions such as these reveals that the requirements on scenarios that are designed to cultivate sound managerial thinking (Lehr et al., 2017) differ markedly from those that are needed to support risk analyses of safety-critical systems such as nuclear waste repositories (Tosoni et al., 2018).

A persistent challenge faced by scenario modelers is the inherent trade-off between specificity and comprehensiveness, or as Kemp-Benedict (2004) calls them "complexity" and "complicatedness". That is, the more details are included in the narrative or the description of uncertainty factors, the more detailed (and possibly also more captivating) the scenario is likely to become; but at the same time, the number of scenarios that would be needed to span the full range of uncertainties would grow dramatically (Carlsen et al., 2016). In the face of this inherent trade-off, most methods of qualitative scenario analysis advocate the formulation of relatively few scenarios that are diverse enough in terms of the outcomes of their uncertainty factors so that they can facilitate, for instance, the formulation of robust decision strategies that perform satisfactorily across all scenarios regardless of what future may bring (Bradfield et al., 2005; Wright and Cairns, 2011). In this case, quantitative methods can be useful, for instance, in choosing those combinations of outcomes that are sufficiently consistent and comprehensive and, as result, lead to the qualitative elaboration of scenarios that are purposely designed to portray the broad span of alternative futures (Ritchey, 2009; Seeve and Vilkkumaa, 2022).

Still, focusing on a small subset of possible scenarios has its limitations. For example, in safety-critical systems, there can be a large number of low-probability high-impact scenarios. These should not be excluded from the analysis, given that none of the scenarios which constitute an unacceptably high risk can be ignored when assessing whether or not the risks are acceptable (Aven and Zio, 2011). Another limitation of focusing only on a small number of qualitatively framed scenarios is that such scenarios cannot be readily integrated with quantitative approaches for providing decision support (e.g. risk or decision analysis). This kind of methodological integration can be achieved only by framing the scenarios within a theoretically rigorous and methodologically coherent framework (Bunn and Salo, 1993). From the viewpoint of defensible decision support, the underpinning uncertainties should be interpreted probabilistically, and sound approaches, as exemplified by the von Neumann-Morgenstern utility theory, should be employed in the assessment of scenario consequences (Rabin, 2013).

By design, methods of cross-impact analysis (CIA) have been crafted for the purpose of examining and characterizing the presence, direction and magnitude of the interdependencies that may exist between the uncertainty factors and their possible outcomes from which scenarios are built. Towards this end, many CIA methods employ ordinal assessment scales for systematically encoding statements about these interdependencies (see, e.g., Scholz and Tietje, 2002; Weimer-Jehle, 2006). In particular, when these judgments are interpreted in terms of probabilities, it is possible to make conclusions about what such interdependencies imply in conjunction with other forms of expert statements (Alter, 1979). For example, the entire set of elicited statements can be synthesized to infer information about the probabilities of specific event outcomes (Gordon and Hayward, 1968; Dalkey, 1971) to identify the most important scenarios (Weimer-Jehle, 2006; Seeve and Vilkkumaa, 2022), or to derive lower and upper bounds the scenario probabilities (Bañuls and Turoff, 2011; Salo et al., 2021).

In this paper, we build on the literature which embraces probability theory as a theoretically coherent framework within which cross-impact statements are interpreted. While the seminal contributions (see, e.g., Gordon and Hayward, 1968; Dalkey, 1971) employ uncertainties with binary outcomes (which either occur or do not), we follow Salo et al. (2021) and view uncertainty factors as random variables which take on one out of several possible outcomes. Specifically, each scenario corresponds to a selection of one outcome for every uncertainty factor. Analyzing probabilities using the methods presented in Salo et al. (2021), however, is computationally onerous, given that the full implications or all elicited judgments are derived by solving a non-linear mathematical optimization model. While this helps preserve the consistency of the elicited judgments, the optimization can be quite time-consuming and inconsistencies in the judgments can be difficult to resolve. As a result, the approach may be impractical if the elicitation process needs to be carried out quickly or when there are no particular requirements to adopt conservatism as a guiding principle in inferring conclusions from the analysis (as is the case in safety-critical systems).

Against this backdrop, our main contribution in this paper is the development of a probabilistic cross-impact method which (i) adopts a well-founded interpretation of cross-impact statements and (ii) derives probability estimates for all possible scenarios (i.e., all combinations of realizations for all uncertainty factors) in a computationally efficient manner during the elicitation process, thereby enabling a host of subsequent analysis. We also demonstrate how the scenario probabilities can be used to construct Bayesian networks, facilitating further analyses.

Furthermore, we describe a case study in which the method was employed to assess the significance of advances in 3D-printing technology for the Finnish Defence Forces. Additive manufacturing, colloquially referred to as 3D-printing, refers to a wide variety of processes that can be used to construct a three-dimensional object based on a digital file (Kietzmann et al., 2015). The methods by which this is achieved include successively depositing, joining, or solidifying relatively thin material layers. In all these methods there is plenty of ongoing research and product development, which accelerates industry growth (Jiménez et al., 2019). In effect, 3D-printing technology shows a lot of promise for military use with applications ranging from entirely novel production methods to spare part logistics (Booth et al., 2018; Heinen and Hoberg, 2019).

The rest of this paper is structured as follows. Section 2 discusses earlier methods of cross-impact analysis. Section 3 formulates our method and its computations. Section 4 outlines the case study. Section 5 concludes.

2 | METHODS OF CROSS-IMPACT ANALYSIS

The origins of cross-impact methods can be traced to the 1960s when Theodore Gordon and Olaf Helmer developed a game called *Future* for the Kaiser Aluminum and Chemical Company (Gordon, 1994). In this game, uncertain events with a given prior probability were written on cards. A die was then rolled to simulate whether or not that event happened. If it did, then the card was flipped over revealing how the probabilities of other events would change as a result. These were the first cross-impacts.

Subsequently, notable contributions to the development and application of cross-impact methods have been made by Gordon (Gordon and Hayward, 1968; Gordon, 1994), Helmer (Helmer, 1977, 1981) and others (Godet, 1976, 1994; Panula-Ontto, 2019). In many of the early methods and their later variants, the probabilities are estimated by considering causal relations between events, even if the temporal occurrence of these events is not necessarily exactly specified. There is the underpinning assumption that the elicitation of conditional probabilities can be instructive in its own right and also potentially cognitively less demanding because in the elicitation of marginal probabilities the respondent would need to take an implicit expectation with regard to all the uncertainty factors whose outcomes are not specified for the event whose probability is being elicited. Computationally, many of these methods can be viewed as computerized implementations of the original card game, in the sense that the event probabilities define Monte Carlo chains in which the cross-impacts cause changes in event probabilities when a different event is realized.

In principle, the sequences of events in these early Monte Carlo simulation methods could be viewed as scenarios. However, these methods are not well suited for estimating the probabilities of all scenarios, because a very large number of simulation runs would be required to obtain accurate results. This would especially be the case for the scenarios with low probabilities which, by definition, would not appear but a small fraction of the total number of simulation runs, but could still give rise to high consequences.

Partly in response to this recognition, dedicated scenario probability estimation methods have been developed. One of the first is presented by Dalkey (1971) who computes a feasible set of scenario probabilities for a consistent cross-impact matrix (i.e a matrix whose elements do not violate the laws of probability theory). Another example is the BASICS tool of Batelle Memorial Institute which computes scenario probabilities using abstracted cross-impact statements instead of strictly probabilistic ones. However, in these methods and in general, the specification of cross-impact estimates that are fully consistent with the tenets of probability theory is not easy for any expert (Huss and Honton, 1987).

The third category of cross-impact methods, which we refer to as structural analysis methods, eschew probabilities altogether. These methods seek to identify key scenarios or uncertainty factors based on the strengths of relationships between the factors as quantified on an ordinal scale. As in BASICS, these scales usually do not have a strictly probabilistic interpretation and they can be fully qualitative. Methods in this category include MICMAC (Godet, 1994), Cross-Impact Balances (Weimer-Jehle, 2006), and the consistency analysis method proposed by Seeve and Vilkkumaa (2022), among others. Because these methods do not involve probabilities they tend to be easy to use and computationally straightforward, making them well-suited for exploratory analyses. However, from a theoretical point of view, without probabilities, they cannot be integrated with probabilistic risk analysis methods or the tenets of expected utility theory.

The method proposed in the present paper is focused on the estimation of scenario probabilities in a setting where some of the cross-impact estimates may be inconsistent. This method is predictable in the sense that the same set of cross-impact estimates always produces the same scenario probabilities, which is in contrast to the presence of some randomness of results obtained by simulation approaches. It also scales up rather well to problems with many uncertainty factors and outcomes, as exemplified by our case study on 3D printing.

3 | METHODOLOGICAL DEVELOPMENT

As in Salo et al. (2021), a *scenario* is here defined as a combination of outcomes of uncertainty factors. The uncertainty factors are modelled as discrete random variables $X^i,\ i = 1, \dots, N$ which have outcomes $S_i = \{1, \dots, n_i\}$. Thus, a *scenario* is a vector $\mathbf{s} = (s_1, \dots, s_N)$ where $s_i \in S_i$ is the outcome for uncertainty factor i. The set of all possible scenarios is the Cartesian product $S := S_{1:N} = \times_{i=1}^{N} S_i$. Thus, the number of all possible scenarios is $|S| = \prod_{i=1}^{N} n_i$.

For a subset of uncertainty factors $F \subseteq \{1, \dots, N\}$, a *partial scenario* is defined as a combination of outcomes for the uncertainty factors which are contained in F. Consequently, the set of partial scenarios for F is $S_F = \times_{i \in F} S_i$. An example of a partial scenario is $\mathbf{s}_{1:i} = (s_1, \dots, s_i)$ which consists of outcomes for the i first uncertainty factors and thus belonging to the set of partial scenarios $S_{1:i} = \times_{j=1}^{i} S_j$. When all uncertainty factors belong to F, then, by construction, the set of corresponding partial scenarios coincides with the set of all scenarios. If only the first i uncertainty factors are considered so that $i < N$, then partial scenarios do not cover outcomes for the uncertainty factors $j > i$.

For the purposes of probabilistic analysis, however, any partial scenario $\mathbf{s}_{1:i} = (s_1, \dots, s_i) \in S_{1:i}$ is compatible with all those (full) scenarios in which the outcomes of the first i uncertainty factors are the same as in the partial scenario $\mathbf{s}_{1:i}$. Thus, any partial scenario $\mathbf{s}_{1:i}$ can be viewed as the collection of those scenarios which can be obtained by extending this partial scenario with outcomes for the uncertainty factors $j = i+1, \dots, N$ so that $E(\mathbf{s}_{1:i}) = \{s' \in S \mid s'_j = s_j,\ \forall j = 1, \dots, i\}$. Furthermore, the probability of the partial scenario $\mathbf{s}_{1:i} \in S_{1:i},\ i \leq N$ can be defined as the sum of the probabilities of those scenarios which can be obtained by extending it to full scenarios so that

$$p(\mathbf{s}_{1:i}) = \sum_{\mathbf{s}_{1:N} \in E(\mathbf{s}_{1:i})} p(\mathbf{s}_{1:N}). \tag{1}$$

Much in the same way, the marginal probability of the outcome $l \in S_j$ for the j-th uncertainty factor is the sum of probabilities for all those scenarios in which this uncertainty factor takes on this outcome, i.e.,

$$P(X^j = l) = \sum_{\mathbf{s} \in S_{1:i} | s_j = l} p(\mathbf{s}). \tag{2}$$

In referring to partial or full scenarios (which correspond to the cases $i < N$ and $i = N$, respectively), we may drop the subscript referring to the number of uncertainty factors. In this case, if $\mathbf{s} \in S_{1:i}$ and $j \in \{1, \dots, i\}$, then s_j refers to the outcome of the j-th uncertainty factor in $\mathbf{s}$.

The above definitions are general in that the uncertainty factors can represent events with binary outcomes (something happens/does not happen) as well as multi-state outcomes (the realization of the i-th uncertainty factor is one of n_i possible outcomes). The setup is broad enough to accommodate real-valued random variables, given that the measurement scale for recording possible outcomes can be typically discretized into a set of disjoint and mutually exhaustive intervals. For example, the rise in global temperatures during the 100-year period from 2000 to 2100 can be categorized as low (< 3°C), medium (3°C–5°C), or high (> 5°C). Because time is a real-valued variable of this kind, it is also possible to include uncertainty factors whose realization indicates when a given event will occur.

In what follows, we develop our method for estimating scenario probabilities in four parts. Section 3.1 presents the basic definitions used for cross-impact multipliers. Section 3.2 describes the basic method for computing the scenario probabilities. Section 3.3 explains how conditional independence information can be harnessed to improve the speed and accuracy of the estimation process. Section 3.4 shows how to build Bayesian networks using conditional independence information and computed probabilities. Finally, Section 3.5 discusses the limitations of the method and its computational properties.

3.1 | Cross-impact multipliers

One rationale for the cross-impact analysis is that the number of scenarios is often so large that it is practically impossible to elicit scenario probabilities directly. For example, from 11 uncertainty factors with 3 possible outcomes each, it is possible to define a total of $3^{11} = 177\,147$ distinct scenarios. In this setup, instead of characterizing the scenario probabilities directly, methods of probabilistic cross-impact analysis characterize probabilistic dependencies between uncertainty factors through cross-impacts and use such characterizations to infer information about the scenario probabilities.

In this paper, we employ the cross-impact interpretation in (Salo et al., 2021), which we call the cross-impact multiplier approach in which the cross-impact between events a and b is defined as

$$C_{ab} := \frac{P(a|b)}{P(a)}, \tag{3}$$

meaning that

$$P(a|b) = C_{ab}P(a). \tag{4}$$

Thus, the cross-impact multiplier specifies how many times more likely the occurrence of the event a becomes when the event b is known to occur. Here, it is worth pointing out that the expressions in (4) do not refer to the temporal sequence in which the events would occur. For instance, it could be the case that the event b will occur after the event a within the time horizon of interest. Still, within such a time horizon, the probability of the event a may be higher when it is known that the event b, too, will come about.

There is some empirical evidence that suggests that humans are more adept at estimating relative magnitudes than providing numerical values (Gallistel and Gelman, 1992). Thus, one motivation for employing the above relative change in probability is that statements about the above ratio (3) may be easier for the experts to estimate than conditional probability values as such.

An appealing property of the cross-impact multiplier is that it is symmetric

$$C_{ab} = \frac{P(a|b)}{P(a)} = \frac{P(a \wedge b)}{P(a)P(b)} = \frac{P(b|a)}{P(b)} = C_{ba}. \tag{5}$$

This reduces the number of cross-impacts to be estimated in half because these multipliers are the same in either direction. Unlike some earlier cross-impact methods, such as the seminal work by Helmer (1981) and Gordon (1994), this interpretation of cross-impacts is not limited to causal relations but reflects also other types of probabilistic dependencies. Specifically, while the presence of a causal relationship does give rise to probabilistic dependence, all probabilistic dependencies between pairs of uncertainty factors cannot be attributed to direct causality between the two uncertainty factors. This would be the case, for instance, when there is a shared underlying cause for two distinct events which are not causally related to each other. To illustrate, consider a situation where two different kinds of alarm systems have been installed for fire detection, one for heat detection and the other for smoke detection. Then neither one of the alarms would cause the other to go off, yet the two alarms would be related to each other in the sense that there would be a positive correlation between them. Analogously, improvements in the affordability and performance of new technologies may exhibit positive correlations even if the underpinning causal determinants of such improvements are not necessarily captured in the model.

In terms of notation, we employ the notation $p(\cdot)$ to refer to the underlying probability distribution $P(\cdot)$ over scenarios. Estimates about scenario probabilities are indicated through $\hat{p}(\cdot)$ where the argument specifies which scenarios are being considered. Thus, for example, for a given outcome $s_i \in S_i$, the $\hat{p}(s_i)$ is the elicited estimate about the marginal probability $p(s_i)$. The probabilities which are derived from the estimates through computations are indicated by $q(\cdot)$. These distinctions are useful in that it is possible, for example, the explore conditions under which the computed probabilities are guaranteed to converge to the true underlying probabilities.

For any $k \in S_i$ and $l \in S_j$, we introduce the abbreviated notations $p^i_k := P(X^i = k)$ and $p^j_l := P(X^j = l)$ for the marginal probabilities and $p^{i|j}_{k|l} = P(X^i = k | X^j = l)$ for the conditional probability. From (4), we get

$$p^{i|j}_{k|l} = C^{ij}_{kl} p^i_k \Leftrightarrow p^{ij}_{kl} = C^{ij}_{kl} p^i_k p^j_l \tag{6}$$

where $p^{ij}_{kl} = P(X^i = k, X^j = l)$ and C^{ij}_{kl} is the cross-impact multiplier for the outcome pair in which the outcome for the uncertainty factor i is k and that for the uncertainty factor j is l.

3.2 | Conditional probability updating

The probabilistic approach proposed by Salo et al. (2021) invites the respondent to specify lower and upper bounds for the cross-impact multiplier (4) and then converts these bounds on the scenario probabilities $p(\mathbf{s}_{1:N}), \forall \mathbf{s}_{1:N} \in S_{1:N}$. Together with estimates of the expected consequences in each of the scenarios, lower and upper bounds for the expected disutility are then derived to provide an aggregate measure of risk. In particular, it is consequently possible to verify whether or not the risk level of the systems is acceptable.

A limitation of this approach is that it presumes that the cross-impact statements elicited from the respondents remain fully consistent (i.e., for any given set of statements that have been elicited from the respondent, there exists at least one assignment of probabilities to all scenarios such that the constraints which correspond to these statements are satisfied). To guide the respondent in providing such statements, however, it is necessary to solve an optimization problem with quadratic constraints which will give rise to computational challenges when the number of uncertainty factors is large. More generally, this approach is not suitable for synthesizing a set of possibly inconsistent cross-impact statements in order to determine a single probability distribution.

Against this backdrop, one of the main contributions of this paper lies in developing a computationally efficient method that (i) admits cross-impacts statements, including inconsistent ones, as well as many other forms of statements that correspond to constraints on scenario probabilities and (ii) synthesizes such statements into a single probability distribution over scenarios in such a way that the resulting distribution represents the best fit to the statements. The estimates from which the scenario probabilities are derived consist of marginal probabilities $\hat{p}^i_k, \hat{p}^j_l$ for all uncertainty factors and their outcomes and cross-impact multipliers $\hat{C}^{ij}_{kl}$ for selected pairs of uncertainty factors and their outcomes. Specifically, the probability distribution over scenarios can be derived even in the absence of information about some pairs of cross-impact multipliers.

To motivate the approach, assume that estimates about the marginal probabilities $\hat{p}^i_k, \hat{p}^j_l$ as well as the cross-impact multiplier $\hat{C}^{ij}_{kl}$ have been elicited. If these estimates are correct in the sense that $\hat{p}^i_k = p^i_k, \hat{p}^j_l = p^j_l$ and $\hat{C}^{ij}_{kl} = C^{ij}_{kl}$, the probability $p^{ij}_{kl} = P(X^i = k, X^j = l)$ is equal to $\hat{C}^{ij}_{kl}\hat{p}^i_k\hat{p}^j_l$. Equivalently, however, this same probability can be expressed as the sum of probabilities for all those scenarios such that $X^i = k$ and $X^j = l$. Thus, we have the following constraint on scenario probabilities

$$\sum_{\substack{\mathbf{s}\in S_{1:N} \\ s_i=k, s_j=l}} p(\mathbf{s}) = \hat{C}^{ij}_{kl}\hat{p}^i_k\hat{p}^j_l, \tag{7}$$

where the summation on the left side is taken over those scenarios whose outcomes for the i-th and j-th uncertainty factors match those on the right side of the equality.

To consider the situation where several (but not necessarily all) cross-impact multipliers have been specified, assume that there exists a binary relation $R_{ij} : S_i \times S_j$ such that $(s_i, s_j) \in R_{ij}$ if and only if the estimate $\hat{C}^{ij}_{kl}$ for the cross-impact multiplier (4) is available. In this case, the above term would appear for all pairs of outcomes such that $R_{ij}(s_i, s_j)$, suggesting that the probability distribution that best matches these estimates can be obtained by solving the minimization problem

$$\min_{p(s)} \sum_{i=2}^{N}\sum_{j=1}^{i-1}\sum_{(k,l)\in R_{ij}} \Big[\Big(\sum_{\substack{\mathbf{s}\in S_{1:N} \\ s_i=k, s_j=l}} p(\mathbf{s})\Big) - \hat{C}^{ij}_{kl}\hat{p}^i_k\hat{p}^j_l\Big]^2 \tag{8}$$

Computationally, however, a concern with the problem (8) is that the optimization would need to be carried out over *all* scenarios. This will be challenging if the number of scenarios is large, either because there are many uncertainty factors or if these factors have several possible outcomes (recall that the total number of scenarios is $\prod_{i=1}^{N} n_i$ where n_i is the number of possible outcomes for the i-th uncertainty factor).

However, the probability of any scenario $s \in S_{1:N}$ can be written as by conditioning the realization of the i-th uncertainty factor on the partial scenario defined by $i-1$ preceding uncertainty factors, i.e.,

$$p(\mathbf{s}) = p(s_i|\mathbf{s}_{1:N\setminus i})p(\mathbf{s}_{1:N\setminus i}) \tag{9}$$

where $\mathbf{s}_{1:N\setminus i}$ is the partial scenario that contains the outcomes for all uncertainty factors except the i-th one. In particular, if the terms $p(\mathbf{s}_{1:N\setminus i})$ representing probabilities for the partial scenarios excluding the i-th uncertainty factor are known, then the estimation of the scenario probabilities becomes a significantly smaller problem in that it is necessary to only consider cross-impact multipliers $\hat{C}^{ij}_{kl}$ that relate to the uncertainty factor i to estimate the conditional probabilities $p(s_i|\mathbf{s}_{1:N\setminus i})$.

Hence, for a given ordering of the uncertainty factors, the relationship (9) can be exploited to build

$$\begin{aligned} p(s) &= p(s_N|\mathbf{s}_{1:N-1})p(\mathbf{s}_{1:N-1}) = p(s_N|\mathbf{s}_{1:N-1})p(s_{N-1}|\mathbf{s}_{1:N-2})p(\mathbf{s}_{1:N-2}) = \dots \\ &= p(s_N|\mathbf{s}_{1:N-1})p(s_{N-1}|\mathbf{s}_{1:N-2})\dots p(s_2|s_1)p(s_1), \end{aligned} \tag{10}$$

This relationship leads to the recognition that the scenario probabilities can be derived iteratively by (i) starting from the marginal probabilities for the first uncertainty factor $p(s_1)$, (ii) computing the conditional probabilities $p(s_2|s_1)$ which represent the best fit to the cross-impact multipliers for the outcomes of the two first uncertainty factors, and (iii) using these conditional probabilities to estimate the probabilities for partial scenarios which comprise these two uncertainty factors. After this step, the iteration can proceed to the third uncertainty factor and so on until all uncertainty factors have been reached.

Thus, the procedure can be described as follows:

1. Use previously computed probabilities for partial scenarios and estimates about marginal probabilities and cross-impact multipliers to compute the conditional probabilities for the next uncertainty factor whose outcomes are conditioned on the previously analyzed partial scenarios.
2. Generate the updated set of partial scenarios which includes this new uncertainty factor. The number of these partial scenarios is equal to the product of (i) the number of partial scenarios in the previous iteration and (ii) the number of outcomes for the new uncertainty factor.
3. Use the computed conditional probability distributions to compute the joint probability distribution for the updated set of partial scenarios.

More formally, the iteration can be carried out by computing probabilities for all partial scenarios $\mathbf{s}_{1:i}, i = 1, \ldots, N$ and with the help of conditional probabilities $q(s_i|\mathbf{s}_{i:i-1})$ such that the iteration is initialized by setting $q(k) \leftarrow \hat{p}^1_k$ for any $k \in S_1 = \{1, \ldots, n_1\}$. At each step of the ensuing iteration, the conditional probabilities can be computed from

$$\min_{q(k|\mathbf{s}_{1:i-1})} \sum_{j=1}^{i-1} \sum_{(k,l)\in R_{ij}} \Big[\Big(\sum_{\{\mathbf{s}\in S_{1:i-1}|s_j=l\}} q(k|\mathbf{s})q(\mathbf{s})\Big) - \hat{C}^{ij}_{kl}\hat{p}^i_k\hat{p}^j_l\Big]^2 \tag{11}$$

$$\sum_{\mathbf{s}\in S_{1:i-1}} q(k|\mathbf{s})q(\mathbf{s}) = \hat{p}^i_k, \qquad \forall\, k \in \{1,2,...,n_i\} \tag{12}$$

$$\sum_{k=1}^{n_i} q(k|\mathbf{s}_{1:i-1}) = 1, \qquad \forall\, \mathbf{s}_{1:i-1} \in S_{1:i-1} \tag{13}$$

$$q(k|\mathbf{s}_{1:i-1}) \geq 0. \qquad \forall\, k \in \{1,2,...,n_i\}, \mathbf{s}_{1:i-1} \in S_{1:i-1} \tag{14}$$

In the third summation of the objective function, the sum is taken over those partial scenarios in which the state of the j-the uncertainty factor is equal to the outcome specified by the term in the relation R_{ij}. The last two constraints ensure that the conditional probability distribution is well-defined. The computed probabilities for the next partial scenarios (which are constructed by appending the states of the i-th uncertainty factor $k \in S_i$ to the previous partial scenarios $\mathbf{s}_{1:i-1}$) can be defined by $q((\mathbf{s}_{1:i-1}, k)) \leftarrow q(\mathbf{s}_{1:i-1})q(k|\mathbf{s}_{1:i-1})$. Thus, the constraint (12) ensures that the marginal probability is the same as the estimated marginal probability $\hat{p}^i_k$ of the outcome $s_i = k$, ensuring that the computed probabilities match the estimated marginal probabilities exactly.

The number of conditional probabilities $q(k|\mathbf{s}_{1:i-1})$ is often significantly higher than the number of cross-impact multipliers which appear in the objective function. This is the case especially when only a fraction of all cross-impact multipliers have been elicited. If estimates about all cross-impact multipliers have been elicited, there are $n_i \sum_{j=1}^{i-1} n_j$ terms in the objective function (11). Equation (12) gives rise to n_i constraints and equation (13) has $\prod_{j=1}^{i-1} n_j$ constraints. The number of parameters is $\prod_{j=1}^{i} n_j$, the same as the number of partial scenarios of length i, and thus the algebraic equation system (12)-(14) is underdetermined so that multiple optimal solutions may exist.

On the other hand, if the objective function (11) is strictly positive, the derived probabilities are not fully consistent with the cross-impact terms. An attractive property of the above optimization formulation is that it is capable of handling situations where the cross-impact statements are not consistent. In this case, at least some of the estimates $\hat{C}^{ij}_{kl}, \hat{p}^i_k, \hat{p}^j_l$ differ from the cross-impact multipliers and marginal probabilities implied by the computed probabilities $q(\mathbf{s}), \mathbf{s} \in S_{1:N}$. The implied cross-impacts $\mathring{C}^{ij}_{kl}$ can be obtained from the computed probabilities as

$$\mathring{C}^{ij}_{kl} = \frac{1}{\hat{p}^i_k\hat{p}^j_l} \sum_{\{\mathbf{s}\in S|s_i=k,s_j=l\}} q(\mathbf{s}). \tag{15}$$

This recognition is useful in that it can be harnessed to support the identification and possible revision of those cross-impact estimates which differ most from the implied cross-impacts, either in absolute terms or in terms of the probabilities for the joint event $X^i = k, X^j = l$ that appears in the objective function (11). Because the marginal probabilities are matched exactly, the cross-impact terms for which the following term is maximized

$$\underset{\substack{i,j \in \{1,\ldots N\} \\ (k,l) \in R_{ij}}}{\text{argmax}} \quad \left| \hat{C}^{ij}_{kl} - \mathring{C}^{ij}_{kl} \right| \tag{16}$$

is the one that deviates most from the implied cross-impact multiplier based on the derived scenario probabilities $q(s)$. On the other hand, the solution

$$\underset{\substack{i,j \in \{1,\ldots N\} \\ (k,l) \in R_{ij}}}{\text{argmax}} \quad \left| \left(\hat{C}^{ij}_{kl} - \mathring{C}^{ij}_{kl} \right) \hat{p}^i_k \hat{p}^j_l \right| \tag{17}$$

helps identify the cross-impact multiplier for which there is the greatest discrepancy between the estimated probability of the event $X^i = k, X^j = l$ and that of the computed probabilities. This analysis can thus be employed to identify and, if need be, revise inconsistent cross-impact multipliers.

Furthermore, the implied cross-impacts can be used to explore which probability distributions other than $q(\mathbf{s})$ would match the given $\hat{C}^{ij}_{kl}, \hat{p}^i_k, \hat{p}^j_l$ equally well. If the cross-impact terms $\mathring{C}^{ij}_{kl}$ that are implied by the computed scenario probabilities are assigned back to (8) instead of $\hat{C}^{ij}_{kl}$, the solution $p(\mathbf{s}) = q(\mathbf{s})$ will make all the sum terms equal to 0, but often $q(\mathbf{s})$ is not unique in this regard. To explore other equally feasible distributions, an optimization problem can be formulated

$$\min_{\mathring{q}(\mathbf{s})} \; f(\mathring{q}) \tag{18}$$

$$\sum_{\substack{\mathbf{s} \in S_{1:N} \\ s_i = k, s_j = l}} \mathring{q}(\mathbf{s}) = \mathring{C}^{ij}_{kl} \hat{p}^i_k \hat{p}^j_l, \qquad \forall\, i \in \{2, \ldots, N\},\; j \in \{1, \ldots, i-1\},\; (k,l) \in R_{ij} \tag{19}$$

$$\sum_{\substack{\mathbf{s} \in S_{1:N} \\ s_i = k}} \mathring{q}(\mathbf{s}) = \hat{p}^i_k, \qquad \forall\, i \in \{1, \ldots, N\},\; k \in \{1, \ldots, n_i\} \tag{20}$$

$$\mathring{q}(\mathbf{s}_{1:N}) \geq 0, \qquad \forall\, \mathbf{s}_{1:N} \in S_{1:N} \tag{21}$$

where $f : \mathbb{R}^{|S|} \to \mathbb{R}$ is chosen to find a scenario probability distribution $\mathring{q}$ with specific properties. To give a few examples, $f(\mathring{q}) = -\mathring{q}(\mathbf{s}^*)$ will maximize the probability of a specific scenario s^* and $f(\mathring{q}) = \sum_{\mathbf{s} \in S} \left(\frac{1}{|S|} - \mathring{q}(\mathbf{s}) \right)^2$ will find the distribution closest to the uniform distribution when $|S|$ is the total number of possible scenarios. Finding expected utility maximizing or minimizing scenario probability distributions is also possible if the utilities of all the scenarios are known. Because all the constraints are linear, the optimization problem can be solved with commonly used optimization tools as long as $f(\mathring{q})$ is convex.

3.3 | Conditional independence

When there are many uncertainty factors, there are inherent limitations in using cross-impact multipliers to estimate all possible scenario probability distributions. To illustrate this, assume that the N uncertainty factors have equally many outcomes n (i.e., $n_i = n,\; i = 1, \ldots, N$). Then the number of different cross-impact multipliers that can be elicited is $\sum_{i=1}^{N} \sum_{j=i+1}^{N} n_i n_j = N(N-1)n^2/2$. This is proportional to the square of the number of uncertainty factors and their outcomes. Still, the number of scenarios $\prod_{i=1}^{N} n_i = n^N$ grows exponentially with the number of uncertainty factors. This implies that estimates about marginal probabilities and cross-impact multipliers will not suffice to fully characterize all possible scenario probabilities, because the number of constraints implied by the cross-impact multipliers will be much lower than the number of scenarios.

Against this backdrop, two observations on the optimization problem in (11)-(14) are in order. First, the number of estimates about cross-impact multipliers increases in every step of the iterative algorithm because the n_i outcomes of the new uncertainty factor are compared with the previously considered uncertainty factors. That is, if all these estimates have been provided, there are $n_i \sum_{j=1}^{i-1} n_j/2$ terms in the objective function (11) while the conditional probabilities $q(k|\mathbf{s}_{1:i-1})$ have already been fixed. As a result, the optimization problems grow in size at every step of the process.

Second, the outcomes of those uncertainty factors which appear earlier on in the sequence of uncertainty factors appear in a larger number of optimization problems. As a result, they likely exert more influence on the final scenario probabilities. More specifically, the probabilities for the partial scenarios defined by the uncertainty factors in the early part of the sequence will not be impacted by the cross-impact terms in the latter part of the sequence. In qualitative terms, this implies that the sequence should be developed so that the uncertainty factors in the early part of the sequence should not be impacted by the later uncertainty factors.

A way to solve this problem of increased complexity is to limit the number of uncertainty factors in every iterative step by focusing only on the relevant dependencies. Specifically, when estimating the conditional probability of the outcome k for the i-th uncertainty factor $p(k|\mathbf{s}_{1:i-1})$, any uncertainty factor whose outcome does not affect this conditional probability is irrelevant. That is, an uncertainty factor a is *irrelevant* for uncertainty factor i in partial scenario set $S_{1:i}$ if and only if $p(k|\mathbf{s}_{1:i-1}) = p(k|\mathbf{s}_{1:i-1\backslash a})$, $\forall k = 1, ..., n_i$, $\forall \mathbf{s}_{1:i-1} \in S_{1:i-1}$, when $\mathbf{s}_{1:i-1\backslash a}$ is the same partial scenario as $\mathbf{s}_{1:i-1}$ but with uncertainty factor a removed. Equivalently, uncertainty factor a is *irrelevant* for i in partial scenario set $S_{1:i}$, if and only if, random variables X_a and X_i are conditionally independent in every partial scenario $\mathbf{s}_{1:i-1\backslash a} \in S_{1:i-1\backslash a}$.

To visualize the conditional dependencies between uncertainty factors, a directed acyclic graph such as the one in Figure 1 can be a useful tool. The uncertainty factors are represented by nodes drawn as circles. The edges connecting the nodes, drawn as arrows, indicate that uncertainty factors are relevant to each other, whereas the lack of a connecting edge implies *irrelevance*. This is how conditional independence between variables is represented in graphical models such as Bayesian networks (Pearl and Paz, 2022).

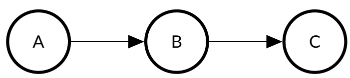

FIGURE 1 *Irrelevance* between uncertainty factors depicted as a graph.

Incorporating conditional independence in the expert judgment elicitation process can be done in two ways. The first is to start by constructing a directed acyclic graph that depicts the dependence structure of the uncertainty factors and then collecting cross-impact information on the directly connected uncertainty factors. The second way is to give the experts the option to state that A does not provide any information about uncertainty factor C that uncertainty factor B does not already provide (as opposed to inviting them to estimate the cross-impacts for outcomes of A and C).

The connection between the graph and the cross-impact matrix containing expert judgments is illustrated in Table

	A	B	C
A		X	
B			X
C			

TABLE 1 When uncertainty factors A and C are conditionally independent, the submatrix describing the cross-impacts between their states can be left empty. Only the submatrices marked with X are filled with cross-impacts.

1. All the submatrices describing the dependencies between outcomes of two uncertainty factors whose outcomes are all either independent or conditionally independent are left empty, whilst the rest are marked with an X. (The cells below the diagonal are not considered, because cross-impacts are symmetric.) The formed matrix can be interpreted as an adjacency matrix for the graph representation, where X in row A column B means that uncertainty factor A has an edge connection to B in the graph. This way, it is possible to either form the graph first and then ask the experts for the cross-impacts only in submatrices marked with X, or the experts can fill out the entire cross-impact matrix freely and leave those submatrices empty that they deem independent or conditionally independent.

Incorporating conditional independence into the computational model is straightforward because multiple *irrelevant* uncertainty factors are also *jointly irrelevant* (i.e. if the two or more uncertainty factors are *irrelevant* any partial scenario formed as a combination of their states is also irrelevant. See A for proof). Let us denote the set of relevant (not *irrelevant*) uncertainty factors for i with D_i and the associated partial scenario set with $S_{D_i} = \bigtimes_{j \in D_i} \{s_1^j, ..., s_{n_j}^j\}$. The probability distribution over partial scenarios in S_{D_i} is calculated from the probability vector p of $S_{1:i-1}$ by taking a sum over all the partial scenarios in $S_{1:i-1}$ that can be obtained by extending the partial scenario $\mathbf{s}_{D_i}$

$$p(\mathbf{s}_{D_i}) = \sum_{\mathbf{s} \in E(\mathbf{s}_{D_i})} p(\mathbf{s}) \tag{22}$$

$$= \sum_{\substack{\mathbf{s}_{1:i-1} \in S_{1:i-1} \\ E(\mathbf{s}_{1:i-1}) \subseteq E(\mathbf{s}_{D_i})}} p(\mathbf{s}_{1:i-1}). \tag{23}$$

Thus, the probability distribution p for partial scenario set S_{D_i} is the marginal distribution for uncertainty factors in D_i. The constraint $E(\mathbf{s}_{1:i-1}) \subseteq E(\mathbf{s}_{D_i})$ means that every scenario in the extension of $\mathbf{s}_{1:i-1}$ can also be found in the extension of $\mathbf{s}_{D_i}$, i.e. partial scenarios $\mathbf{s}_{1:i-1}$ and $\mathbf{s}_{D_i}$ have the same outcomes for all uncertainty factors in D_i.

Now, because the uncertainty factors that are not included in D_i are *irrelevant* for estimating the conditional probabilities for the i-th uncertainty factor based on the partial scenarios $\mathbf{s}_{1:i-1}$, we have

$$p(s_k^i | \mathbf{s}_{1:i-1}) = p(s_k^i | \mathbf{s}_{D_i}), \quad \text{if} \quad E(\mathbf{s}_{1:i-1}) \subseteq E(s_{D_i}). \qquad \forall\ \mathbf{s}_{1:i-1} \in S_{1:i-1},\ \mathbf{s}_{D_i} \in S_{D_i} \tag{24}$$

Furthermore, when cross-impact statements are available for pairs of uncertainty factors i, j such that $j < i$ and $j \notin D_i$, the optimization problem (11-14) can be solved in S_{D_i} instead of $S_{1:i-1}$ so that

$$\min_{q(k|\mathbf{s}_{D_i})} \sum_{j\in D_i} \sum_{(k,l)\in R_{ij}} \Big[\Big(\sum_{\{\mathbf{s}\in S_{D_i}|s_j=l\}} q(k|\mathbf{s})q(\mathbf{s})\Big) - \hat{C}^{ij}_{kl}\hat{p}^i_k\hat{p}^j_l\Big]^2 \tag{25}$$

$$\sum_{\mathbf{s}\in S_{D_i}} q(k|\mathbf{s})q(\mathbf{s}) = \hat{p}^i_k, \qquad \forall\, k \in \{1,2,\dots,n_i\} \tag{26}$$

$$\sum_{k=1}^{n_i} q(k|\mathbf{s}_{D_i}) = 1, \qquad \forall\, \mathbf{s}_{D_i} \in S_{D_i} \tag{27}$$

$$q(k|\mathbf{s}_{D_i}) \geq 0, \qquad \forall\, k \in \{1,2,\dots,n_i\},\ \forall\, \mathbf{s}_{D_i} \in S_{D_i} \tag{28}$$

and then (24) can be used to get probabilities for all partial scenarios in $S_{1:i-1}$. This makes it easier to solve the problem because the number of scenarios in the optimization will now depend only on the relevant uncertainty factors D_i instead of all factors $\{1,\dots,i-1\}$. The inclusion of conditional independence also leads to more tightly constrained scenarios without adding any additional optimization constraints.

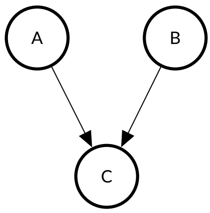

FIGURE 2 Two independent uncertainty factors A and B which share a dependent uncertainty factor C.

Implementing conditional independence, however, adds some additional limitations on the order in which the uncertainty factors are included in the iterative process, because the uncertainty factors upon which the conditional independence relies are already contained in the partial scenario set $S_{1:i-1}$. Similarly, uncertainty factors that are marginally independent, such as A and B in Figure 2, but have other uncertainty factors dependent on them, should be included before the dependencies. This is because while introducing additional *irrelevant* uncertainty factors in expanding the set of partial scenarios cannot turn *irrelevant* uncertainty factors into relevant ones, the same does not hold for adding relevant uncertainty factors, which can under specific circumstances change other uncertainty factors from *irrelevant* into relevant.

3.4 | Bayesian networks

Conditional probability information and the conditional independence structure can also be combined to form a Bayesian network (Pearl and Paz, 2022). In the network, the conditional independence information is represented by a directed acyclic graph, like the one in Figure 3. The nodes (circles) represent uncertainty factors and the edges (arrows) between nodes represent conditional dependencies. Because the graph is directed and acyclic, there exists at least one total ordering of nodes such that node u precedes node v if there is an edge from node u to node v in the graph. This is called a *topological ordering* of the graph. In Figure 3 alphabetical order is one possible *topological ordering*. If two nodes do not share an edge they are conditionally independent given a subset of other nodes that precede either of the two. Conversely, if two nodes share an edge, they are not conditionally independent.

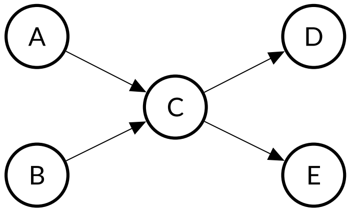

FIGURE 3 A simple conditional independence network.

Probability information is incorporated in the Bayesian network in the form of conditional probability tables like the one seen in Figure 4. Assuming the indexing of uncertainty factors follows a *topological ordering* of the Bayesian network, the probability table of uncertainty factor i contains the conditional probabilities $p(k|\mathbf{s}_{D_i}) \ \forall k \in \{1, ..., n_i\}, \ \mathbf{s}_{D_i} \in S_{D_i}$. When combined, the conditional probability distributions $p(k|\mathbf{s}_{D_i})$ of all the uncertainty factors can be used to calculate the probability of any scenario $p(\mathbf{s})$ in S. This is because of the chain rule of probability (10) and conditional independence (24)

$$\begin{aligned} p(s) &= p(s_N|s_{1:N-1})p(\mathbf{s}_{1:N-1}) = p(s_N|s_{D_N})p(\mathbf{s}_{1:N-1}) \\ &= p(s_N|s_{D_N})p(s_{N-1}|s_{1:N-2})p(s_{1:N-2}) = p(s_N|s_{D_N})p(s_{N-1}|s_{D_{N-1}})p(\mathbf{s}_{1:N-2}) \dots \\ &= p(s_N|s_{D_N})p(s_{N-1}|s_{D_{N-1}}) \dots p(s_2|s_{D_2})p(s_1), \end{aligned} \tag{29}$$

which makes these conditional probability distributions a very memory-efficient way of storing the probability distribution.

Constructing the Bayesian network based on cross-impacts and conditional independence information is quite straightforward because the computational estimates for $p(k|\mathbf{s}_{D_i})$ can be obtained by solving the optimization problem (25-28). Indexing the uncertainty factors following a *topological ordering* of the conditional dependence graph is not a problem, because the cross-impact multipliers only measure probabilistic dependence and not causality, and thus they do not limit the order in any way.

3.5 | Computational considerations

The results of the iterative method may differ slightly from those obtained by solving a single optimization problem to determine all scenario probabilities which would represent the best fit to all the elicited expert judgments about marginal probabilities and cross-impacts. Specifically, in the notation of Section 3.2, the problem of fitting scenario probabilities directly can be stated as the linear least squares problem

$$\min_{q(\mathbf{s})} \sum_{i=2}^{N} \sum_{j=1}^{i-1} \sum_{(k,l)\in R_{ij}} \Big[\Big(\sum_{\substack{\mathbf{s}\in S_{1:N} \\ s_i=k, s_j=l}} q(\mathbf{s})\Big) - \hat{C}_{kl}^{ij} \hat{p}_k^i \hat{p}_l^j \Big]^2 \tag{30}$$

$$\sum_{\substack{\mathbf{s}\in S_{1:N} \\ s_i=k}} q(\mathbf{s}) = \hat{p}_k^i, \qquad \forall\, i \in \{1, ..., N\},\ k \in \{1, ..., n_i\} \tag{31}$$

$$0 \le q(\mathbf{s}) \le 1, \qquad \forall\, \mathbf{s} \in S_{1:N}. \tag{32}$$

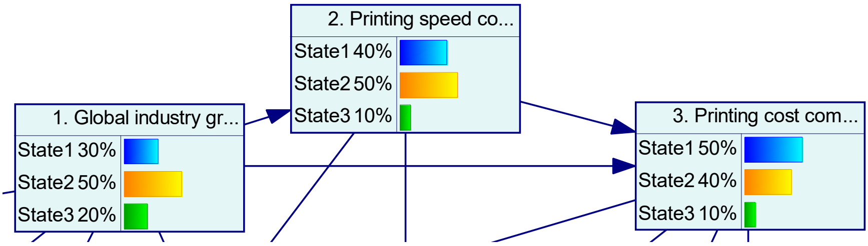

		Global industry growth								
		s_1^1			s_2^1			s_3^1		
		Printing speed								
		s_1^2	s_2^2	s_3^2	s_1^2	s_2^2	s_3^2	s_1^2	s_2^2	s_3^2
	$\mathbf{s}_{D_3}$	(s_1^1, s_1^2)	(s_1^1, s_2^2)	(s_1^1, s_3^2)	(s_2^1, s_1^2)	(s_2^1, s_2^2)	(s_2^1, s_3^2)	(s_3^1, s_1^2)	(s_3^1, s_2^2)	(s_3^1, s_3^2)
Printing cost	$p(s_1^3 \mid \mathbf{s}_{D_3})$	0.88	0.81	0.65	0.55	0.31	0.33	0.31	0.09	0.40
	$p(s_2^3 \mid \mathbf{s}_{D_3})$	0.12	0.19	0.35	0.40	0.68	0.33	0.34	0.09	0.34
	$p(s_3^3 \mid \mathbf{s}_{D_3})$	0	0	0	0.05	0.01	0.33	0.35	0.82	0.26

FIGURE 4 One conditional probability table from the Bayesian network in the case study. The rows of the table show the outcome probabilities of the third uncertainty factor, Printing Cost, conditioned on the relevant preceding uncertainty factor's outcomes.

where the next to last constraint ensures that the scenario probabilities are matched to the marginal probabilities. The difference between the iterative process and this direct fitting stems from the fact that the iterative method does not weigh all cross-impact statements simultaneously. However, our computational experiments suggest that when the statements about marginal probabilities and cross-impacts are consistent (i.e., the sum (30) is zero for some scenario probabilities $q(s)$), both approaches produce similar distributions, which also fit all the marginals and cross-impacts perfectly by construction.

However, while the formulation (30) is conceptually simpler than the iterative approach in Section 3.2, it has a major caveat in that it leads to much bigger optimization problems. Specifically, the size of the optimization problem in (30) grows exponentially and in our computational tests workstations with 16-32GM RAM ran out of memory when the number of scenarios reached tens of thousands. In particular, the case study in Section 4 with its $3^{11} = 177\,147$ scenarios proved too large for this approach. In contrast, the iterative approach which exploits judgments about conditional independence was able to construct the scenario probability distribution for the case study in 3.9 seconds using a laptop with 2.40GHz I5 processor and 16GB RAM.

In short, the iterative method performs better when analyzing large systems with more uncertainty factors. The reason for this is that the size of the optimization problems does not depend on the total number of scenarios but, rather, only on the size of the partial scenario set which contains the relevant uncertainty factors. To illustrate this point, if the new uncertainty factors can be added without increasing the average number of relevant factors per each new factor, the computational complexity of the iterative process increases linearly with the number of total uncertainty factors. This can be contrasted with the direct fitting approach in which the size of the optimization problems grows exponentially so that these problems become quickly unsolvable. In other words, while direct scenario probability fitting can still be used in smaller problems, the iterative method will prove indispensable in many problems of realistic size.

4 | CASE STUDY

We next present a case study on analyzing the developing 3D-printing technologies and their future impact on the Finnish Defence Forces (FDF). The aim of this case study was to (i) identify uncertainty factors that have a large impact on, how 3D printing would be applied in the Finnish military in view of possible developments over the next 15 years, (ii) specify ranges of possible realizations for these uncertainty factors, (iii) characterize dependencies between the uncertainty factors, and (iv) build structured scenario framework which would capture these multiple inputs, by doing so, and (v) serve as a tool for offering insights into questions which are of focal concern to the FDF in the context of 3D printing.

4.1 | Uncertainty factors

The identification of uncertainty factors was preceded by a systematic literature review and preliminary interviews with experts, resulting in an initial set of ten key uncertainty factors with three outcomes for each. These uncertainty factors were discussed at length in a 4-hour remote workshop which was organized by using video conferencing tools and attended by a panel consisting of half eight 3D-printing experts from the Finnish military and research community. The specific fields of expertise represented by the panelists covered military logistics, 3D-printing business, and 3D-printing technology.

In the workshop, the experts reached a consensus that the factor *Progress in 3D manufacturing* should be separated into two factors, representing *Printing speed* and *Printing costs*. This lead to the final list of uncertainty factors in Table 2.

In the workshop, the outcomes of every uncertainty factor were discussed together with the experts. A verbal description was developed for each, including numerical bounds where appropriate. For each outcome of every uncertainty factor, the corresponding marginal probability was assigned. This represented the baseline probability of this outcome in the absence of information about the outcomes of other uncertainty factors.

Next, the experts were asked to characterize cross-impacts using a seven-point scale from -3 to 3. This scale was employed to record statements about how the probability of a given outcome for an uncertainty factor would change from its baseline probability as a result of knowing that another uncertainty factor will have a specific outcome. For example, how much more likely it is that the global 3D industry maintains the growth speed of 2019-2020 if the costs associated with printing fall by 50-90%? A small part of the cross-impact matrix is in Table 3.

The statements recorded on this ordinal scale were converted into estimates about cross-impact multipliers through the transformation

$$C_{kl}^{ij} = \sqrt{2}^{V_{kl}^{ij}}, \tag{33}$$

where C_{kl}^{ij} is the cross-impact multiplier derived from the statement V_{kl}^{ij}. Thus, responses from the range -3 to 3 were mapped to numerical values cross-impact multipliers $\frac{1}{3},\ \frac{1}{2},\ \frac{2}{3},\ 1,\ 1\frac{1}{2},\ 2$, and 3, and this information was available to the experts during the evaluations. The experts only estimated the cross-impacts of those uncertainty factor pairs

Uncertainty Factor	Outcome	Probability
1. Global industry growth	Decrases	0.3
	Remains same	0.5
	Increases	0.2
2. Printing speed compared to present	Up to 2 times faster	0.4
	2 to 10 times faster	0.5
	Over 10 times faster	0.1
3. Printing cost compared to present	Up to 50% cheaper	0.5
	50% to 90% cheaper	0.4
	Over 90% cheaper	0.1
4. Finnish industry growth	Decreases	0.3
	Remains same	0.5
	Increases	0.2
5. Graduates with 3D-printing expertise	Up to 100 (current)	0.2
	100 to 300	0.6
	Over 300	0.2
6. Legal regulation of 3D-printing in Finland	Limits strongly	0.05
	Similar to other manufacturing	0.9
	No regulation	0.05
7. Standardization of processes and models	No standardization	0.35
	Includes technical requirements	0.45
	Full automation possible	0.2
8. Use of 3D-printed objects in FDF	Just individual items	0.1
	Common and has purchase procedures	0.5
	Access to 3D-printing capacity on demand	0.4
9. FDF access to 3D-printing model files	Just individual items	0.2
	Relevant models included in system purchases	0.7
	Models available for most new and old systems	0.1
10. FDF 3D-printing spare parts in peacetime	Low importance	0.7
	Significant and well planned	0.29
	Crucial and strictly controlled	0.01
11. FDF 3D-printing spare parts in crisis times	Low importance	0.45
	Significant and well planned	0.45
	Crucial and strictly controlled	0.1

TABLE 2 Uncertainty factors, their outcomes, and corresponding marginal probabilities estimated for the year 2035.

they deemed to provide significant information about each other. The remaining pairs were deemed conditionally independent given the preceding uncertainty factors, as is shown in Table 4.

			Printing cost		
			50-100%	10-50%	Max 10%
			0.5	0.4	0.1
Global industry growth	Slow	0.3	2	0	-2
	Stable	0.5	0	2	0
	Fast	0.2	-2	0	3
Printing speed	100-200%	0.4	1	-1	-1
	200-1000%	0.5	-1	1	0
	Over 1000%	0.1	-1	0	1

TABLE 3 Part of the cross-impact estimate matrix on probabilistic dependencies between pairs of outcomes for uncertainty factors.

	1	2	3	4	5	6	7	8	9	10	11
1. Global industry growth		X	X	X	X	X	X				
2. Printing speed compared to present			X	X							X
3. Printing cost compared to present				X				X	X	X	
4. Finnish industry growth					X	X	X	X	X	X	X
5. Number of graduates with 3D-printing expertise									X		
6. Legal regulation of 3D-printing in Finland								X	X	X	
7. Standardization of processes and models									X		X
8. Use of 3D-printed objects in FDF									X	X	X
9. FDF access to 3D-printing model files										X	X
10. FDF 3D-printing spare parts in peacetime											X
11. FDF 3D-printing spare parts in crisis times											

TABLE 4 The X:s denote the uncertainty factors whose cross-impacts were evaluated. The empty white cells are conditionally independent uncertainty factor pairs.

The usual scale from -3 to 3 was chosen (instead of asking about the cross-impact multipliers directly) to expedite the elicitation process while recognizing that the resulting estimates would not necessarily be consistent answers. Indeed, there were some inconsistencies in the resulting estimates. In Table 3, the highlighted cross-impact between stable industry growth and 10-50% printing cost leads to a situation where the conditional probabilities associated with the row for global industry growth would not sum up to one. This inconsistency is easy to spot, because all terms on this row are non-negative, meaning that stable industry growth would invariably preserve or increase the probabilities of all outcomes of printing costs; but this is impossible because these outcomes are (meant to be) mutually exhaustive (Salo et al., 2021). Similar, and also less apparent inconsistencies appear all over the cross-impact matrix. Indeed, the fact that such inconsistencies in expert judgments are likely to surface in the cross-impact analysis is one of the reasons which motivated us to develop a method that could derive scenario probabilities even when the estimates are not perfectly consistent. This can make the elicitation process both faster and less arduous.

4.2 | Results

We used the presented iterative method to compute the scenario probability distribution for all the scenarios that can be formed from the uncertainty factors in Table 2. The entire distribution could not be included here, because it has $3^{11} = 177\,147$ probabilities. Thus, we are offering some (hopefully) interesting observations instead.

The cross-impact judgments provided by the experts indicated that the role of 3D-printing in the future of spare parts logistics is quite uncertain. There is a 41% probability that it will not have a great role in either the peace or crisis time logistics. From the scenario probability distribution, we calculated that the probability of any scenario where spare parts production in either peace or crisis time is significant or crucial is practically zero, if the use of 3D printed objects in FDF is limited to just individual items or access to 3D-printing models is extremely limited. Collecting a library of 3D-printing model files and building processes to order and use 3D-printed items takes a significant amount of time and effort, so it would be advisable to start as soon as possible if the 3D printing of spare parts is seen as worth pursuing.

Using the conditional probability distributions and conditional independence information (Table 4), we also constructed a Bayesian network using the GeNIe Modeler software (BayesFusion, 2021), seen in Figure 5. The uncertainty factors were introduced starting with exogenous factors that would not be affected by choices the FDF makes, followed by exogenous factors that could be affected in limited ways in cooperation with Finnish government entities and industry, and the last factors included were endogenous to the Finnish military. Thus, factors 1-3 describe the state of the 3D-printing industry globally, 4-7 describe the situation in Finland, and 8-11 describe the situation inside the FDF. The constructed network can be used to illuminate various what-if (partial) scenarios (Fenton and Neil, 2001).

To give a concrete example, looking at the partial scenarios consisting of uncertainty factors 1-7, i.e. the exogenous factors, the most probable is the one in which every factor gets the second outcome. Its probability is 8.62%, which is quite high considering these uncertainty factors can produce $3^7 = 2\,187$ different partial scenarios. The probabilities of the FDF endogenous uncertainty factors in this partial scenario can be seen in Figure 6. 3D-printed parts are very likely to have at least significant importance in crisis time operations. At the same time, they are quite unlikely to be that important during peacetime. Because crisis capabilities are developed during peacetime, this means that special attention should be focused on both training and developing processes to support this 3D-printing spare parts in a crisis, because it seems unlikely to develop on its own.

Another relatively high probability (2.28%) exogenous partial scenario in which uncertainty factors 1-5 and 7 all obtain the first outcome while uncertainty factor 6 obtains its second outcome 2 is in Figure 7. This is a more pessimistic partial scenario for the 3D-printing industry as a whole and represents growth and technological development slowing down significantly. Here 3D-printing is unlikely to play any role at all in the spare parts logistics, and this helps showcase why there is such a high probability of them remaining unremarkable despite their seeming importance in the most probable scenario in Figure 6. The developments in the industry as a whole are going to have an impact on

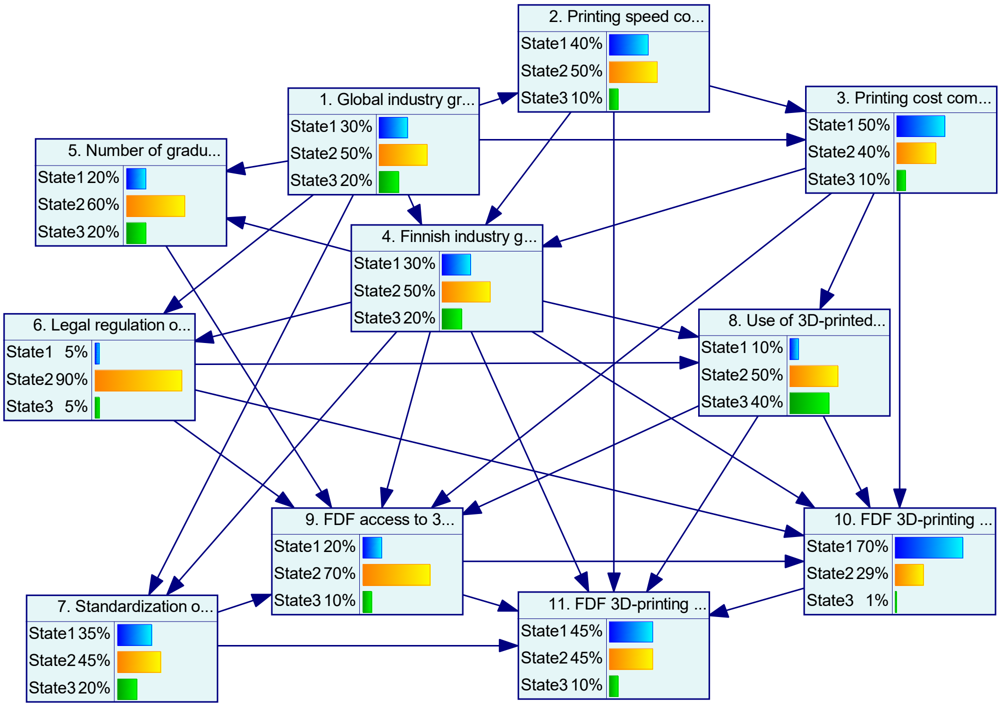

FIGURE 5 The constructed Bayesian network in GeNIe Modeler software (BayesFusion, 2021).

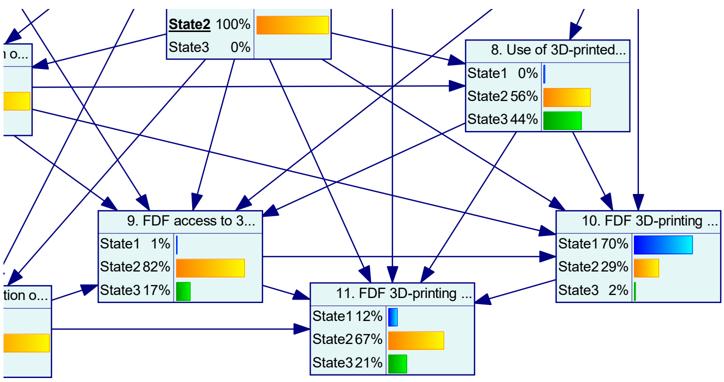

FIGURE 6 The probability distributions of uncertainty factors describing 3D-printing in Finnish Defence Forces in the most likely exogenous partial scenario.

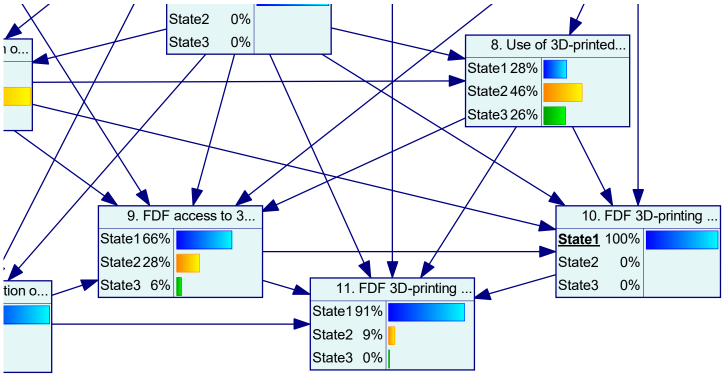

FIGURE 7 The probability distributions of uncertainty factors describing 3D-printing in Finnish Defence Forces in a pessimistic high probability exogenous partial scenario.

the usefulness of the technology for the FDF and should be monitored carefully.

5 | CONCLUSIONS AND DISCUSSION

Scenario analysis provides a structured framework for identifying and exploring key uncertainties which shape the future. Because the range of qualitative and quantitative scenario methods is so wide, there are different perspectives on the rationales and the suitable ways of carrying out scenario analysis (Millett, 2009), but qualitative and quantitative scenario methods are not fundamentally at odds. Rather, they are complementary: for instance, evocative narratives can be made even more compelling by accompanying them with numerical data while detailed quantitative analyses can be enriched with storylines to communicate the implications behind the numbers. In short, the choice of methodologies should be guided by how the scenarios are going to be used.

Against this backdrop, we have formulated a method to elicit expert judgments about cross-impact terms which are processed to infer the accompanying joint probability distribution over all possible scenarios. A notable benefit of this approach is that it facilitates the integration with other well-founded quantitative approaches – including expected utility theory, probabilistic risk assessment, and statistical inference – and thus expands the extant range of available techniques for foresight and strategic planning.

Scenario probabilities and Bayesian networks depicting dependencies between the uncertainty factors have many practical uses in the military context. Scenario probability distributions facilitate a number of different analyses to assess the impacts of new technologies. Numerous simulation (Lappi, 2008; Rao et al., 1993), game theoretic (Poropudas and Virtanen, 2010; Roponen et al., 2020) and dynamic (Gue, 2003) tools can be used to analyze the scenario-specific system performance, but their ability to support strategic analysis is limited without the underlying scenario probabilities.

The same also applies to technology forecasting beyond the military context to some extent. For example, climate models help generate well-founded scientific predictions concerning the rate of change in the global temperature and sea level, rise, but they are less apt at predicting what kinds of mitigation actions governments will take or how people will respond to changing environmental conditions; yet such behavioral would also need to be accounted for in order to address risks comprehensively. As a result, there is a need for scenario models such as ours which harnesses cross-impact statements in order to link technological changes to the key behavioral responses that are pivotal in shaping the future.

Among cross-impact methods, ours is purposely grounded on the estimation of all possible scenario probabilities. Most probabilistic cross-impact methods tend to rely on Monte Carlo simulation, which, however, tends to require an impractically large number of iterations to reach good accuracy when the number of scenarios is large. Computationally, our method scales well into problems with even dozens of uncertainty factors, especially if the number of probabilistic dependencies between the uncertainty factors is not too large. The problems caused by large dependency sets can be mitigated to some extent by choosing the included uncertainty factors and iteration order in the right way, but eventually, a limit is reached on how much can be expressed with just pairwise dependency statements. This is a limitation shared by all cross-impact techniques because the number of possible scenarios grows faster than the number of cross-impacts.

We have employed unconditional cross-impact multipliers (4) in which the relative change in the probability of a given outcome level does not explicate assumptions about the realizations of uncertainty factors beyond the two that are considered in the comparison. Mathematically, however, one could elicit conditional cross-impact multipliers which would explicate such assumptions with no reason why the cross-impact multipliers could not be even extended into triplets or a larger number of uncertainty factors. We have chosen not to explore it beyond conditional independence in this paper, because the number of triplets and beyond grow so much faster than the number of pairs, that collecting such information for all uncertainty factors would be practically infeasible in most cases. However, introducing individual optimization constraints based on higher-level dependencies would be straightforward if desired.

Although our case study has focused on 3D-printing, the proposed method is generic and can be readily applied across numerous contexts in which it is of interest to build a comprehensive model that retains all possible scenarios. Thus, its advantages lie in countering the risk that the focus is, perhaps prematurely, placed on a small subset of scenarios, as opposed to capturing the full breadth of possible scenarios that can be built as combinations of outcomes of several uncertainty factors. In such contexts, cross-impact analysis offers a pragmatic, relatively straightforward, and cognitively manageable approach to assessing dependencies between uncertainty factors. Furthermore, the models proposed in this paper are computationally efficient and make it possible to provide informative insights based on the interactive exploration of the implications of all model inputs, including judgments about the marginal outcome probabilities and cross-impact statements.

Acknowledgements

The experts who participated in our workshop and provided their valuable insights for the case study are gratefully acknowledged. This research has been financially supported by MATINE - The Scientific Advisory Board for Defence of the Finnish Ministry of Defence.

Conflict of interest

The authors declare that there are no conflicts of interest.

References

Alter, S. (1979) The evaluation of generic cross-impact models. *Futures*, **11**, 132–150.

Aven, T. and Zio, E. (2011) Some considerations on the treatment of uncertainties in risk assessment for practical decision making. *Reliability Engineering & System Safety*, **96**, 64–74.

Bañuls, V. A. and Turoff, M. (2011) Scenario construction via Delphi and cross-impact analysis. *Technological Forecasting and Social Change*, **78**, 1579–1602.

BayesFusion, L. (2021) *GeNIe Modeler*. Available online: https://www.bayesfusion.com/ (accessed on 31 May 2022).

Booth, J., Edwards, E., Whitley, M., Kranz, M., Seif, M. and Ruffin, P. (2018) Military comparison of 3d printed vs commercial components. In *Nano-, Bio-, Info-Tech Sensors, and 3D Systems II*, vol. 10597, 52–66. SPIE.

Bowman, G. (2016) The practice of scenario planning: An analysis of inter-and intra-organizational strategizing. *British Journal of Management*, **27**, 77–96.

Bradfield, R., Wright, G., Burt, G., Cairns, G. and Van Der Heijden, K. (2005) The origins and evolution of scenario techniques in long range business planning. *Futures*, **37**, 795–812.

Bunn, D. W. and Salo, A. A. (1993) Forecasting with scenarios. *European Journal of Operational Research*, **68**, 291–303.

Carlsen, H., Eriksson, E. A., Dreborg, K. H., Johansson, B. and Bodin, Ö. (2016) Systematic exploration of scenario spaces. *Foresight*, **18**, 59–75.

Chermack, T. J. (2022) *Using Scenarios: Scenario Planning for Improving Organizations*. Berrett-Koehler Publishers.

Dalkey, N. C. (1971) An elementary cross-impact model. *Technological Forecasting and Social Change*, **3**, 341–351.

Duinker, P. N. and Greig, L. A. (2007) Scenario analysis in environmental impact assessment: Improving explorations of the future. *Environmental Impact Assessment Review*, **27**, 206–219.

Fenton, N. and Neil, M. (2001) Making decisions: Using Bayesian nets and MCDA. *Knowledge-Based Systems*, **14**, 307–325.

Gallistel, C. R. and Gelman, R. (1992) Preverbal and verbal counting and computation. *Cognition*, **44**, 43–74.

Godet, M. (1976) Scenarios of air transport development to 1990 by SMIC 74—a new cross-impact method. *Technological Forecasting and Social Change*, **9**, 279–288.

— (1986) Introduction to la prospective: Seven key ideas and one scenario method. *Futures*, **18**, 134–157.

— (1994) *From Anticipation to Action: A Handbook of Strategic Prospective*. UNESCO publishing.

Gordon, T. J. (1994) *Cross-impact method*, vol. 4. American Council for the United Nations University.

Gordon, T. J. and Hayward, H. (1968) Initial experiments with the cross impact matrix method of forecasting. *Futures*, **1**, 100–116.

Gue, K. R. (2003) A dynamic distribution model for combat logistics. *Computers & Operations Research*, **30**, 367–381.

Heinen, J. J. and Hoberg, K. (2019) Assessing the potential of additive manufacturing for the provision of spare parts. *Journal of Operations Management*, **65**, 810–826.

Helmer, O. (1977) Problems in futures research: Delphi and causal cross-impact analysis. *Futures*, **9**, 17–31.

— (1981) Reassessment of cross-impact analysis. *Futures*, **13**, 389–400.

Huss, W. R. and Honton, E. J. (1987) Scenario planning—what style should you use? *Long Range Planning*, **20**, 21–29.

Jiménez, M., Romero, L., Domínguez, I. A., Espinosa, M. d. M. and Domínguez, M. (2019) Additive manufacturing technologies: An overview about 3d printing methods and future prospects. *Complexity*, **2019**.

Kemp-Benedict, E. (2004) From narrative to number: A role for quantitative models in scenario analysis. In *2nd International Congress on Environmental Modelling and Software*.

Kietzmann, J., Pitt, L. and Berthon, P. (2015) Disruptions, decisions, and destinations: Enter the age of 3-D printing and additive manufacturing. *Business Horizons*, **58**, 209–215.

Kosow, H. and Gaßner, R. (2008) *Methods of Future and Scenario Analysis: Overview, Assessment, and Selection Criteria*, vol. 39. DEU.

Lappi, E. (2008) Sandis military operation analysis tool. In *2nd Nordic Military Analysis Symposium*, 17–18.

Lehr, T., Lorenz, U., Willert, M. and Rohrbeck, R. (2017) Scenario-based strategizing: Advancing the applicability in strategists' teams. *Technological Forecasting and Social Change*, **124**, 214–224.

Millett, S. M. (2009) Should probabilities be used with scenarios. *Journal of Futures Studies*, **13**, 61–68.

O'Neill, B. C., Carter, T. R., Ebi, K., Harrison, P. A., Kemp-Benedict, E., Kok, K., Kriegler, E., Preston, B. L., Riahi, K., Sillmann, J. et al. (2020) Achievements and needs for the climate change scenario framework. *Nature Climate Change*, **10**, 1074–1084.

Panula-Ontto, J. (2019) The axiom approach for probabilistic and causal modeling with expert elicited inputs. *Technological Forecasting and Social Change*, **138**, 292–308.

Pearl, J. and Paz, A. (2022) Reasoning about relevance relations. *Probabilistic and Causal Inference: The Works of Judea Pearl*, 189.

Pereira, H. M., Leadley, P. W., Proença, V., Alkemade, R., Scharlemann, J. P., Fernandez-Manjarrés, J. F., Araújo, M. B., Balvanera, P., Biggs, R., Cheung, W. W. et al. (2010) Scenarios for global biodiversity in the 21st century. *Science*, **330**, 1496–1501.

Poropudas, J. and Virtanen, K. (2010) Game-theoretic validation and analysis of air combat simulation models. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, **40**, 1057–1070.

Rabin, M. (2013) Risk aversion and expected-utility theory: A calibration theorem. In *Handbook of the Fundamentals of Financial Decision Making: Part I*, 241–252. World Scientific.

Rao, A. S., Lucas, A., Morley, D., Selvestrel, M. and Murray, G. (1993) Agent-oriented architecture for air-combat simulation. *Tech. rep.*, Technical Note 42, The Australian Artificial Intelligence Institute.

Ritchey, T. (2009) Developing scenario laboratories with computer-aided morphological analysis. In *14th International Command and Control Research and Technology Symposium, Washington DC*.

Roponen, J., Insua, D. R. and Salo, A. (2020) Adversarial risk analysis under partial information. *European Journal of Operational Research*, **287**, 306–316.

Salo, A., Tosoni, E., Roponen, J. and Bunn, D. W. (2021) Using cross-impact analysis for probabilistic risk assessment. *Futures & Foresight Science*, **4**, e2103.

Scholz, R. W. and Tietje, O. (2002) *Embedded Case Study Methods: Integrating Quantitative and Qualitative Knowledge*. Sage.

Schwartz, P. (2012) *The Art of the Long View: Planning for the Future in an Uncertain World*. Currency.

Seeve, T. and Vilkkumaa, E. (2022) Identifying and visualizing a diverse set of plausible scenarios for strategic planning. *European Journal of Operational Research*, **298**, 596–610.

Siljander, R. and Ekholm, T. (2018) Integrated scenario modelling of energy, greenhouse gas emissions and forestry. *Mitigation and Adaptation Strategies for Global Change*, **23**, 783–802.

Tosoni, E., Salo, A. and Zio, E. (2018) Scenario analysis for the safety assessment of nuclear waste repositories: A critical review. *Risk Analysis*, **38**, 755–776.

Weimer-Jehle, W. (2006) Cross-impact balances: A system-theoretical approach to cross-impact analysis. *Technological Forecasting and Social Change*, **73**, 334–361.

Wright, G. and Cairns, G. (2011) *Scenario Thinking: Practical Approaches to the Future*. Palgrave Macmillan London.

A | PROOFS

Definition An uncertainty factor a is *irrelevant* for uncertainty factor i in scenario set $S_{1:i}$ if and only if $p(s^i_k|\mathbf{s}_{1:i-1}) = p(s^i_k|\mathbf{s}_{1:i-1\backslash a})$, $\forall k = 1,...,n_i$, $\forall \mathbf{s}_{1:i-1} \in S_{1:i-1}$, when $\mathbf{s}_{1:i-1\backslash a}$ is the same scenario as $\mathbf{s}_{1:i-1}$ but without the outcome of uncertainty factor a.

Remark Definition A comes with the built-in assumption that probabilities $p(\mathbf{s}_{1:i-1})$ are strictly positive because otherwise, the conditional probabilities would be undefined. They can be arbitrarily small, however.

Theorem 1 (Intersection) *If uncertainty factors a and b are irrelevant for uncertainty factor i in scenario set $S_{1:i}$, then $p(s^i_k|\mathbf{s}_{1:i-1}) = p(s^i_k|\mathbf{s}_{1:i-1\backslash a,b})$, $\forall k = 1,...,n_i$, $\forall \mathbf{s}_{1:i-1} \in S_{1:i-1}$.*

Proof We assume that a, b and i are uncertainty factors in {1,...,i} and $a \neq b, b \neq i, i \neq a$. Assuming uncertainty factors a and b are separately *irrelevant* for uncertainty factor i in scenario set $S_{1:i}$, we show that $p(s^i_k|\mathbf{s}_{1:i-1}) = p(s^i_k|\mathbf{s}_{1:i-1\backslash a,b})$, $\forall k = 1,...,n_i$, $\forall \mathbf{s}_{1:i-1} \in S_{1:i-1}$, when $\mathbf{s}_{1:i-1\backslash a,b}$ is the same scenario as $\mathbf{s}_{1:i-1}$ but without the outcome of uncertainty factors a and b, i.e. any partial scenario formed from outcomes of a and b is also irrelevant.

From the assumptions we get,

$$p(s^i_k|\mathbf{s}_{1:i-1}) = p(s^i_k|\mathbf{s}_{1:i-1\backslash a}), \quad \forall k = 1,...,n_i, \quad \forall \mathbf{s}_{1:i-1} \in S_{1:i-1} \tag{34}$$

$$p(s^i_k|\mathbf{s}_{1:i-1}) = p(s^i_k|\mathbf{s}_{1:i-1\backslash b}), \quad \forall k = 1,...,n_i, \quad \forall \mathbf{s}_{1:i-1} \in S_{1:i-1} \tag{35}$$

and thus

$$p(s^i_k|\mathbf{s}_{1:i-1\backslash a}) = p(s^i_k|\mathbf{s}_{1:i-1\backslash b}), \quad \forall k = 1,...,n_i, \quad \forall \mathbf{s}_{1:i-1} \in S_{1:i-1}. \tag{36}$$

Now, let $k \in \{1,...,n_i\}$ and $\mathbf{s}_{1:i-1} \in S_{1:i-1}$. Using the law of total probability, we get

$$p(s^i_k|\mathbf{s}_{1:i-1\backslash a,b}) = \sum_{l=1}^{n_a} p(s^i_k|\mathbf{s}_{1:i-1\backslash a,b}, s^a_l)p(s^a_l|\mathbf{s}_{1:i-1\backslash a,b}) \tag{37}$$

Because $(\mathbf{s}_{1:i-1\backslash a,b} \wedge s^a_l)$ is a scenario in $S_{1:i-1\backslash b}$, we can use the equality from (36), and get

$$p(s^i_k|\mathbf{s}_{1:i-1\backslash a,b}) = \sum_{l=1}^{n_a} p(s^i_k|\mathbf{s}_{1:i-1\backslash a})p(s^a_l|\mathbf{s}_{1:i-1\backslash a,b}) \tag{38}$$

$$= p(s^i_k|\mathbf{s}_{1:i-1\backslash a}) \sum_{l=1}^{n_a} p(s^a_l|\mathbf{s}_{1:i-1\backslash a,b}) \tag{39}$$

$$= p(s^i_k|\mathbf{s}_{1:i-1\backslash a}). \tag{40}$$

Thus, because $p(s^i_k|\mathbf{s}_{1:i-1\backslash a,b}) = p(s^i_k|\mathbf{s}_{1:i-1\backslash a})$ and $p(s^i_k|\mathbf{s}_{1:i-1}) = p(s^i_k|\mathbf{s}_{1:i-1\backslash a})$, it follows that

$$p(s^i_k|\mathbf{s}_{1:i-1}) = p(s^i_k|\mathbf{s}_{1:i-1\backslash a,b}), \qquad \forall k = 1,...,n_i, \quad \forall \mathbf{s}_{1:i-1} \in S_{1:i-1}. \tag{41}$$

It would also be straightforward to extend this proof to any number of *irrelevant* uncertainty factors.

Aalto-DT 67/2023

ISBN 978-952-64-1260-3 (printed)
ISBN 978-952-64-1261-0 (pdf)
ISSN 1799-4934 (printed)
ISSN 1799-4942 (pdf)

Aalto University
School of Science
Department of Mathematics and Systems Analysis
www.aalto.fi

BUSINESS +
ECONOMY

ART +
DESIGN +
ARCHITECTURE

SCIENCE +
TECHNOLOGY

CROSSOVER

DOCTORAL
THESES