INNOCENZO GASPARINI INSTITUTE FOR ECONOMIC RESEARCH

Institutional Members: CEPR, NBER and Università Bocconi

# WORKING PAPER SERIES

## Risk Analysis and Decision Theory: Foundations

*E. Borgonovo, V. Cappelli, F. Maccheroni, M. Marinacci*

**Working Paper n. 556**

**This Version: July 31, 2015**

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
http://www.igier.unibocconi.it

# Risk Analysis and Decision Theory: Foundations

E. Borgonovo, V. Cappelli, F. Maccheroni, M. Marinacci

Department of Decision Sciences and IGIER, Università Bocconi, Milano, Italy.

July 31, 2015

**Abstract**

The triplet-based risk analysis of Kaplan and Garrick (1981) is the keystone of state-of-the-art probabilistic risk assesment in several applied fields. This paper performs a sharp embedding of the elements of this framework into the one of formal decision theory, which is mainly concerned with the methodological and modelling issues of rational decision making. In order to show the applicability of such an embedding, we also explicitly develop it within a nuclear probabilistic risk assessment, as prescribed by the U.S. NRC.

The aim of this exercise is twofold: on the one hand, it gives risk analysis a direct access to the rich toolbox that decision theory has developed, in the last decades, in order to deal with complex layers of uncertainty; on the other, it exposes decision theory to the challenges of risk analysis, thus providing it with broader scope and new stimuli.

*50% of the problems in the world result [...] from people using different words with the same meaning.*

Kaplan (1997, p. 408)

## 1 Introduction

In the management of complex technological systems, the term risk analysis refers to the part of the policy-making process associated with the identification of scenarios and their likelihoods

(Clemen and Reilly, 1999 and Paté-Cornell and Dillon, 2006). The location of a nuclear waste repository (Garrick and Kaplan, 1999), the programming of a space mission (Stamatelatos et al., 2011, Borgonovo and Smith, 2011, Dillon et al., 2003), and the evaluation of design changes in chemical and nuclear plants (Boykin et al., 1984, Caruso et al., 1999), are a few examples in which decision making is informed by a risk analysis, in the so-called *risk-informed decision making* (Apostolakis, 2004). This discipline has *gained a significant amount of attention from both policymakers and the public over the past 30 years, as the interaction of technology and policy choices has become more predominant in the evaluation of trade-offs in a democratic society* (Apostolakis and Pickett, 1998, p. 621).[1]

Over the years, Kaplan and Garrick's definition of risk (Kaplan and Garrick, 1981, henceforth KG) has become one of the pillars of risk analysis, guiding several key studies performed by national and international agencies and laboratories (for instance, Kaplan and Garrick's risk triplets are a structural part of NASA's recent risk management handbook by Stamatelatos et al., 2011). The triplet structure introduced by Kaplan and Garrick remains influential also for recent generalizations of the risk concept (Althaus, 2005, Aven, 2012b). These latter works signal a common trait of risk analysis, that is, the consideration of risk as a self-standing concept, independent of any decision-analytical consideration. This separation is deemed attractive by some researchers (see the debates reported in Aven, 2012), insofar it permits the extension of Kaplan and Garrick's definition of risk to non-probabilistic approaches. However, it has the drawback of neglecting the operative motivation that a decision-making approach brings to a risk analysis.[2]

Indeed, in their seminal 1981 article, Kaplan and Garrick maintain that *risk must thus be considered always within a decision theory context* (KG, p. 25). Several subsequent works discuss risk analysis from a decision-making viewpoint (Howard, 1988, Apostolakis, 1990, Winkler, 1996, and Apostolakis and Pickett, 1998). Both Paté-Cornell and Dillon (2006) and Clemen and Reilly (1999) underline that risk analysis and decision analysis are intertwined: *a decision analysis can include a risk analysis component* (Paté-Cornell and Dillon, 2006, p. 220). Nonetheless, the decision-analytical background upholding Kaplan and Garrick's definition itself has not been investigated in depth to date.

This gap in risk analysis leads us to the decision side. Subjective expected utility originates in the seminal works of von Neumann and Morgenstern (1947), Wald (1950), and Savage (1954), which have become the pillars of modern decision analysis (Pratt et al., 1995, and Smith and von Winterfeldt, 2004). This theory features a decision maker (DM) who evaluates acts whose consequences depend on states of the environment generated by mechanisms that are only partially known or understood. Each such mechanism corresponds to a probabilistic model that describes the frequency of the various states inherent to the phenomenon at hand. The information available to the DM allows him to posit a set of possible mechanisms, that is, of possible probabilistic models.[3] In general, such a set is not a singleton because information is not sufficiently accurate to pin down a single mechanism. In other words, the DM is uncertain about the true probabilistic model. Thus, two layers of uncertainty are at play as follows: the irreducible *aleatory* uncertainty (physical risk) about states and the *epistemic* uncertainty about models (of the physical phenomenon). The first is described by probabilistic models, the second by a prior probability over them. The recent Cerreia-Vioglio et al. (2013), henceforth CMMM, extends Savage's analysis by showing that: if the DM's preferences satisfy Savage (1954)'s axioms plus a consistency condition, then one obtains a subjective expected utility functional where the qualitative distinction of the two layers is meaningful (see Section 5).

The heart of the present work is the reconciliation of Kaplan and Garrick's definition of risk in its various formats with the corresponding decision-theoretic counterparts. First of all, we observe that the notion of hazard in Kaplan and Garrick's is in one-to-one correspondence with the decision-theoretic notion of act. We then show that Kaplan and Garrick's risk triplets in "frequency" format can be embedded in the von Neumann and Morgenstern decision-theoretic framework, in which uncertainty is described by an objective probability. Similarly, triplets in "probability" format belong to Savage's framework, in which uncertainty is described by a subjective probability. Finally, we show that the "probability of frequency" format, where the two layer distinction between aleatory and epistemic uncertainty is achieved, finds its natural collocation within the Waldean extension of Savage's framework recently proposed by CMMM.

The correspondence we obtain is "non-utilitarian," in the sense that it does not rely on the specification of a utility function over outcomes. The purpose of building this "bridge" is twofold:

- it gives risk analysis a direct access to the rich toolbox that decision theory has developed in the last decades in order to deal with complex layers of uncertainty;
- it exposes decision theory to the challenges of risk analysis, thus giving it broader scope and new stimuli.

The relevance of the latter point is self-explanatory. As to the first point, it is important to recall that, starting with the seminal paper of Schmeidler (1989) and until the present days, decision theory has been studying in careful detail the problem of ambiguity which is very closely related to that of epistemic uncertainty (Cerreia-Vioglio et al., 2013b, henceforth CMMMb). Thus the present paper opens the way towards more sophisticated approaches to risk-informed decision making in the presence of epistemic uncertainty, and it consolidates the existing ones (Section 6). Another advantage of connecting risk analysis and decision theory is that the decision theoretic approach is axiomatic; that is, it makes explicit the choice implications of the different criteria available to the DM. Clearly, the unveiling of such implications allows the regulator to better evaluate the criteria adopted by a DM (Apostolakis, 2014).

We illustrate the applicability of our exercise through several examples. In particular, we embed the three levels of nuclear Probabilistic Risk Assessment (PRA) in a decision-theoretic setup.

The remainder of this paper is organized as follows. In Sections 2 and 3 we present, respectively, the KG triplets' setup and the decision theoretic framework. In Section 4, the heart of the paper, we embed the former setup in the latter. This embedding allows us, in Section 5, to connect standard decision making criteria in risk analysis with classical expected utility theories, and so to root these criteria in the axiomatic approach to rational choice, that is proper of decision theory. Conversely, in Section 6, by building on recent developments of decision theory,[4] we propose some robust decision making criteria as risk informed approaches that take into account the qualitative differences between epistemic and aleatory uncertainty (emphasized in risk analysis by, e.g., Paté-Cornell and Fischbeck, 1992, Der Kiureghian and Ditlevsen, 2009, and Marzocchi, Newhall and Woo, 2012). The different uses of those criteria in a risk analysis perspective are discussed in the subsequent Section 7. Finally, Section 8 illustrates the obtained identification in the context of PRA, as prescribed by the U.S. NRC, and Section 9 concludes.

## 2 The Risk Analysis Setup

In their definition of risk, KG state that a risk analysis must be capable of responding to the following three questions:

1. "What can happen?"

2. "How likely is it that it will happen?"

3. "If it does happen, what are the consequences?"

The response to the first question identifies a *scenario* $S_n$. The answer to the second question indicates the *likelihood* $\ell_n$ of scenario $S_n$. The answer to the third question is the consequence $x_n$ of $S_n$. Typically, $x_n$ is a measure of the damage that the occurrence of scenario $S_n$ causes.

*Hazard* is then formally defined by KG (p. 13, footnote 2) as the set of doublets

$$H = \{\langle S_n, x_n \rangle : n = 1, ..., N\} \tag{1}$$

*Risk* is defined as the set of triplets

$$R = \{\langle S_n, \ell_n, x_n \rangle : n = 1, ..., N\}. \tag{2}$$

A term $\langle S_n, \ell_n, x_n \rangle$ in eq. (2) is called a risk triplet. Informally, KG describe the concepts of *hazard*, as "a source of danger", and *risk* as "the possibility of loss or injury" and "the degree of probability of such loss."

Kaplan and Garrick then consider three different formats in which the notion of likelihood takes alternative meanings.

**Format 1 (Frequency)** Consider a repetitive situation underlaid by a stationary and ergodic process with many past observations that allow the DM to learn frequencies. One can then ask "how frequently does scenario $n$ occur?." The likelihood of $S_n$ is expressed by an *objective frequency*:

$$\ell_n = \phi_n.$$

Then, the *frequency format* of risk (2) becomes

$$R = \{\langle S_n, \phi_n, x_n \rangle : n = 1, ..., N\} . \tag{3}$$

If consequences are scalars monotonically ranked as $0 \leq x_1 \leq \cdots \leq x_N$, risk $R$ determines a staircase function $\Phi : [0, x_N] \to [0, 1]$, called *risk curve*, defined by

$$\Phi(x) = \begin{cases} \sum_{n=1}^{N} \phi_n & x \in [0, x_1] \\ \sum_{n=2}^{N} \phi_n & x \in (x_1, x_2] \\ \cdots & \cdots \\ \phi_{N-1} + \phi_N & x \in (x_{N-2}, x_{N-1}] \\ \phi_N & x \in (x_{N-1}, x_N] \end{cases}$$

The value $\Phi(x)$ above is the cumulative probability of a consequence equal to or greater than $x$.

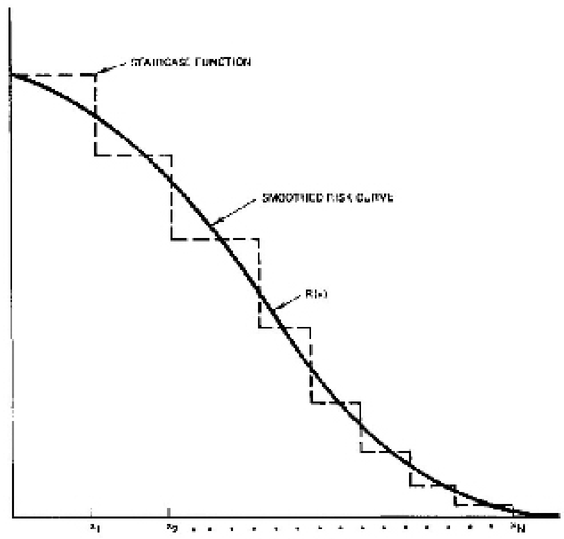


Fig. 1. Risk curve.

Figure 1. A risk curve as in KG

Figure 1 represents a risk curve. When consequences are measured on a continuous scale, using the original expression of KG, the staircase function (dotted) is replaced by the *smoothed* risk curve.

**Format 2 (Probability)** In this format, Kaplan and Garrick consider a general uncertain situation, possibly not repetitive. Here likelihoods quantify the DM's beliefs about scenarios. The likelihood of $S_n$ is thus expressed by a *subjective probability*:

$$\ell_n = p_n.$$

The risk defined by (2) becomes

$$R = \{\langle S_n, p_n, x_n \rangle : n = 1, ..., N\}. \tag{4}$$

Kaplan (1997) calls (4) *risk in probability format* and contrasts it with the frequency format.

In case of scalar consequences, with $0 \leq x_1 \leq \cdots \leq x_N$, risk $R$ here determines a *subjective risk curve.*

**Format 3 (Probability of Frequency)** Finally, consider a repetitive (at least conceptually) situation in which the DM can conceive frequencies, but without enough information to assess them precisely. Then, there is uncertainty about the frequency of scenario $S_n$, which is described by Kaplan and Garrick through a discrete probability distribution:

$$\ell_n = \{\langle \rho_\theta, (\phi_\theta)_n \rangle : \theta \in \Theta\} \tag{5}$$

where $\rho_\theta$ is the subjective probability that $(\phi_\theta)_n$ is the actual frequency of $S_n$. The probability distribution in (5) is called the *probability of frequency* of $S_n$.

With the abbreviation $\ell_n = \rho_n (\phi_n)$ of (5), the risk triplets in (2) give rise to

$$R = \{\langle S_n, \rho_n (\phi_n), x_n \rangle : n = 1, ..., N\} \tag{6}$$

which is therefore called *risk in probability of frequency format.*

<u>Remarks</u>

First, note that $\ell_n = \rho_n (\phi_n)$ is a probability distribution, while in the previous two formats $\ell_n = \phi_n$ and $\ell_n = p_n$ are scalars.

Second, while $\phi_n$ is an objective probability and $p_n$ is a subjective probability, $\rho_n (\phi_n)$ is an hybrid consisting of two parts: $\phi_\theta$ is an objective probability, that captures the aleatory uncertainty about the scenario $S_n$ that will obtain, $\rho_\theta$ is a subjective probability, that captures the epistemic uncertainty about the true $\phi_\theta$. The coexistence of frequentist and Bayesian components is motivated by the copresence of two different sources of risk ("state risk" and "model risk").

Third, it is worth pointing out that $R$ in (6) bundles together the family $\{R_\theta : \theta \in \Theta\}$ of risks in frequency format

$$R_\theta = \{\langle S_n, (\phi_\theta)_n, x_n \rangle : n = 1, ..., N\}$$

by means of the discrete probability distribution $\rho = \{\langle \rho_\theta, \phi_\theta \rangle : \theta \in \Theta\}$. ▲

In the case of scalar consequences, with $0 \leq x_1 \leq \cdots \leq x_N$, the family of risks $\{R_\theta : \theta \in \Theta\}$ determines a family $\{\Phi_\theta : \theta \in \Theta\}$ of risk curves indexed by $\theta$. Formally,

$$\Phi_\theta(x) = \begin{cases} \sum_{n=1}^{N} (\phi_\theta)_n & x \in [0, x_1] \\ \sum_{n=2}^{N} (\phi_\theta)_n & x \in (x_1, x_2] \\ \dots & \dots \\ (\phi_\theta)_{N-1} + (\phi_\theta)_N & x \in (x_{N-2}, x_{N-1}] \\ (\phi_\theta)_N & x \in (x_{N-1}, x_N]. \end{cases}$$

While, graphically, Figure 2 displays the (smoothed) families of risk curves obtained as a result of the risk assessment of a nuclear waste isolation study (Paté-Cornell, 1999).

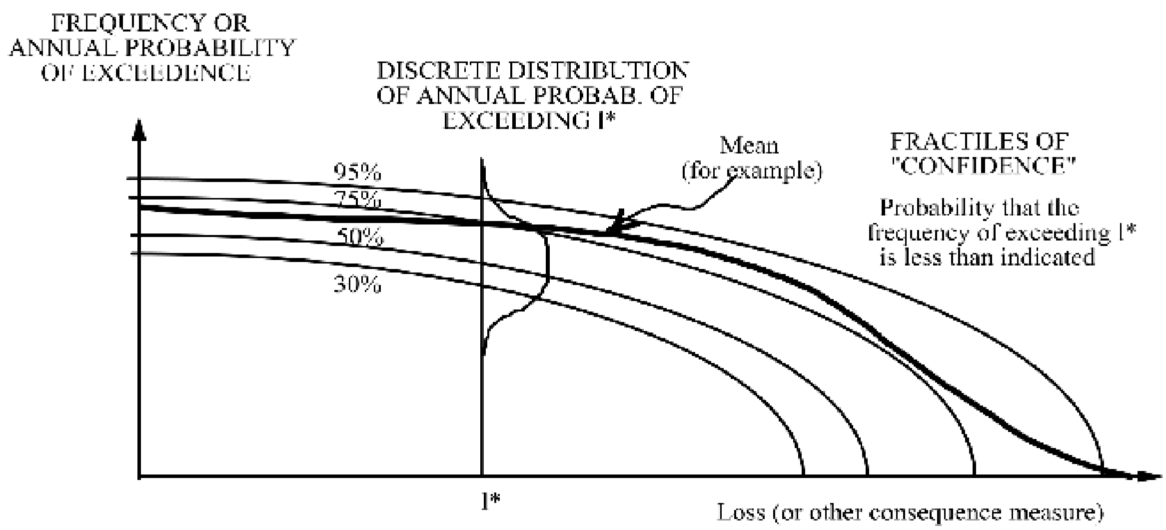


Figure 2. A family of risk curves as in Pate-Cornell (1999).

After enumerating the previous formats, Kaplan (1997, p. 409) writes that, among them, the probability of frequency format "*is the most general and by far the most powerful and useful idea.*" In fact, consider a risk $R = \{\langle S_n, \rho_n(\phi_n), x_n\rangle : n = 1, ..., N\}$ in probability of frequency format. The expected frequency

$$\bar{p}_n = \mathbb{E}[\rho_n(\phi_n)] = \sum_{\theta \in \Theta} (\phi_\theta)_n \rho_\theta$$

can be seen as the probability assigned to scenario $S_n$ (KG, p. 19). This leads to a reduced risk $\bar{R} = \{\langle S_n, \bar{p}_n, x_n \rangle : n = 1, ..., N\}$ in probability format.[5] On the other hand, $\rho_n(\phi_n)$ reduces to $\phi_n$ when only one frequency $\phi$ is considered, that is, $\Theta$ is a singleton. If so, $R$ has the frequency format $\{\langle S_n, \phi_n, x_n \rangle : n = 1, ..., N\}$.

Summing up, as anticipated by KG, the *probability of frequency* format encompasses both the *probability* and the *frequency* formats.

# 3 Decision theoretic setup

## 3.1 The Savage setup

Following Savage (1954), a decision problem under uncertainty features a DM who has to choose among a set of alternative acts whose consequences depend on uncertain factors beyond his control, called states. The consequence of act $f$ in state $s$ is denoted by $f(s)$.

In decision theory, it is common use to denote the set of consequences by $X$ and the state space by $S$. The state space is endowed with an event $\sigma$-algebra $\Sigma$.[6] An act is a simple $\Sigma$-measurable function

$$\begin{aligned} f: \; S & \;\to\; X \\ s & \;\mapsto\; x = f(s) \end{aligned}$$

that is

$$f(s) = \begin{cases} x_1 & s \in S_1 = f^{-1}(x_1) \\ x_2 & s \in S_2 = f^{-1}(x_2) \\ \dots & \dots \\ x_N & s \in S_N = f^{-1}(x_N) \end{cases}$$

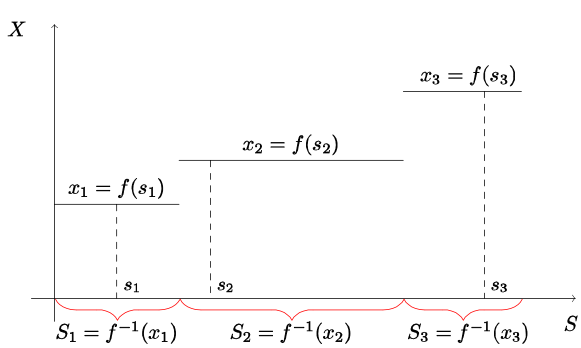


Figure 3. A decision theoretic act

where $\{x_1, x_2, ..., x_N\} = f(S)$ is a finite set of consequences in $X$ and $\{S_1, S_2, ..., S_N\} = \left\{f^{-1}(x_1), ..., f^{-1}(x_N)\right\}$ is a partition of $S$ in $\Sigma$. This observation allows us to write an act of decision theory as an hazard of risk analysis. Namely,

$$f \equiv \left\{\left\langle f^{-1}(x), x\right\rangle : x \in f(S)\right\} = \{\langle S_i, x_i\rangle : n = 1, ..., N\} = H$$

with $S_i = f^{-1}(x_i)$ or, equivalently, $x_i = f(S_i)$.[7] This modelling stage is both "non-utilitarian," in the sense of the introduction, and "pre-probabilistic" in that no assumption whatsoever has still been made about the likelihoods of events.

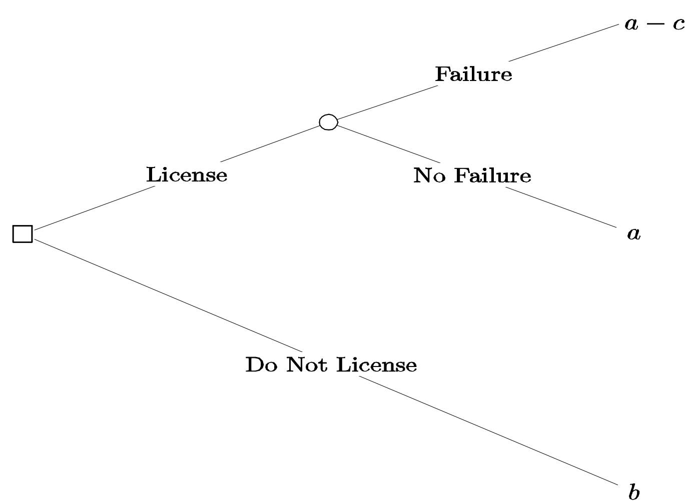


Figure 4. A decision tree as in Howard (1988)

**Example 1** *A Probabilistic Risk Assessment study might be conducted to support the licensing of a new nuclear power plant. Howard (1988) discusses the underlying decision problem through a simple but effective example (Figure 4). The available alternatives are: to licence $L$ or not to licence $NL$. If the plant is licensed and operates properly (that is, it does not fail within the period of interest), it provides a net monetary benefit $a$ to the society. Should the plant fail, society incurs in a large cost $c$, to be subtracted from the benefit. If the plant, instead, is not licensed, society is provided with some benefit $b$. Suppose $\mathbf{X}$ is the vector of random variables that determines the failure conditions of the plant and $g$ is its limit state function. The implied failure and safe event* *are* $S_{\text{failure}} = [g(\mathbf{X}) < 0]$ *and* $S_{\text{safe}} = [g(\mathbf{X}) \geq 0]$, *respectively. The Savage acts corresponding to* $L$ *and to* $NL$ *are:*

$$f_L(s) = \begin{cases} a - c & s \in S_{\text{failure}} \\ \\ a & s \in S_{\text{safe}} \end{cases} \qquad \textit{and} \qquad f_{NL}(s) = b \quad s \in S.$$

*In this case,*

$$S_{\text{failure}} = f_L^{-1}(a-c), \qquad S_{\text{safe}} = f_L^{-1}(a), \qquad \textit{and} \qquad S = f_{NL}^{-1}(b)$$

*so that*

$$f_L \equiv \{\langle S_{\text{failure}}, a-c\rangle, \langle S_{\text{safe}}, a\rangle\} = H_L \qquad \textit{and} \qquad f_{NL} \equiv \{\langle S, b\rangle\} = H_{NL}$$

*where* $H_L$ *and* $H_{NL}$ *are the hazards connected to licensing and not licensing the plant.* ▲

### 3.2 A Waldean extension

In several situations, the DM might not possess enough information to posit a single objective probability (frequency) over the states. He assumes that states $s$ are generated by a random mechanism $m$ that belongs to a given finite collection $\mathcal{M}$ of probabilistic models. Technically, $\mathcal{M}$ is a set of probability measures on $\Sigma$. Each such $m$ represents aleatory uncertainty, that is, the inherent randomness that states feature. In other words, the DM posits a model space $\mathcal{M}$ in addition to the state space $S$.

Denoting by $\mathcal{A}$ the subset of available acts (in Example 1 above, $\mathcal{A} = \{f_L, f_{NL}\}$), the quartet

$$(\mathcal{A}, S, X, \mathcal{M})$$

represents the Savage-Wald form of a (statistical) decision problem (see CMMM and CMMMb).

According to the Bayesian paradigm, the subjective state-of-knowledge of the DM about models in $\mathcal{M}$ is represented by a prior probability $\mu$ on $\mathcal{M}$. The prior $\mu$ on models naturally induces a predictive probability $P_\mu$ on states through model averaging. In decision theory, this process is called *reduction* and the average of conditional models is called *predictive probability*, formally defined by

$$P_\mu(E) = \sum_{m \in \mathcal{M}} m(E)\,\mu(m) \qquad E \in \Sigma.$$

The decision theoretic relation between $\mu$ and $P_\mu$ is discussed in CMMM.

# 4 Embedding

Let $(\mathcal{A}, S, X, \mathcal{M})$ be a decision problem and $\mu$ be a prior probability on $\mathcal{M}$. In order to facilitate the comparison with the risk analysis notation, write

$$\mathcal{M} = \{m_\theta : \theta \in \Theta\}.$$

**Acts and hazards** We already established the relation

$$f \equiv \left\{\left\langle f^{-1}(x), x\right\rangle : x \in f(S)\right\} = \{\langle S_i, x_i\rangle : n = 1, ..., N\} = H \tag{7}$$

between Savage's acts $f$ and Kaplan and Garrick's hazards $H$.

Remark This also implies that, the fourth standard risk analysis question (Greenberg et al., 2012):

4. "How can concequences be prevented or reduced?"

can be rephrased as "what is the set $\mathcal{A}$ of feasible acts?". This set will be typically determined by the hazard containment regulations, the technical limits, and the economic constraints of a society.

▲

**Models and frequencies** For every probabilistic model $m_\theta \in \mathcal{M}$ and every scenario $S_n$, the frequency $(\phi_\theta)_n$ of $S_n$ is $m_\theta(S_n)$, briefly

$$m_\theta \equiv \phi_\theta \tag{8}$$

for all $\theta \in \Theta$. That is, each probabilistic model $m_\theta \in \mathcal{M}$, in the decision theoretic framework, corresponds to a possible frequency $\phi_\theta$ in the Kaplan and Garrick probability of frequency setup.

**Prior and probability of frequency** Having established that each possible frequency $\phi_\theta$ of scenarios corresponds to a probabilistic model $m_\theta$, the probability of frequency $\phi_\theta \equiv m_\theta$ be the true one is simply $\mu(m_\theta)$, that is,

$$\rho_\theta \equiv \mu(m_\theta) \qquad \theta \in \Theta. \tag{9}$$

But then, for every scenario $S_n$,

$$\rho_n(\phi_n) = \{\langle \rho_\theta, (\phi_\theta)_n \rangle : \theta \in \Theta\} = \{\langle \mu(m_\theta), m_\theta(S_n) \rangle : \theta \in \Theta\}$$

and it corresponds to the discrete subjective distribution of the probability evaluation map $m_\theta \mapsto m_\theta(S_n)$, that is,

$$\{\langle \mu(m_\theta), m_\theta(S_n) \rangle : \theta \in \Theta\} \equiv \sum_{\theta \in \Theta} \mu(m_\theta)\, \delta_{m_\theta(S_n)}$$

and the latter is denoted by $\mu^{S_n}$. We thus have

$$\rho_n(\phi_n) \equiv \mu^{S_n}. \tag{10}$$

**Risks** Just putting together all the collected pieces, we obtain

$$\begin{aligned} R &= \{\langle S_n, \rho_n(\phi_n), x_n \rangle : n = 1, ..., N\} \\ &= \left\{\left\langle S_n, \mu^{S_n}, x_n \right\rangle : n = 1, ..., N\right\} \\ &= \left\{\left\langle f^{-1}(x), \mu^{f^{-1}(x)}, x \right\rangle : x \in f(S)\right\} \equiv (f \mid \mu) \end{aligned}$$

where $f$ is the act corresponding to hazard $\{\langle S_i, x_i \rangle : n = 1, ..., N\}$ and $\mu$ is the prior corresponding to $\rho$ as in (9).

**Risk curves** In the case of scalar consequences, with $0 \leq x_1 \leq \cdots \leq x_N$, the family of risk curves $\{\Phi_\theta : \theta \in \Theta\}$ corresponds to the family $\left\{G_f^\theta : \theta \in \Theta\right\}$ of decreasing distribution functions

$$G_f^\theta(x) = m_\theta[s \in S : f(s) \geq x] \qquad x \in [0, x_N] \tag{11}$$

of act $f$ under the models $m_\theta$. In turn, each $G_f^\theta$ can be identified with the its discrete distribution

$$\left(m_\theta \circ f^{-1}\right)(x) = m_\theta[s \in S : f(s) = x] \qquad x \in f(S).$$

These distributions are called lotteries in decision theory.[8] That is, risk curves are, essentially, the lotteries induced by the act corresponding to the faced hazard, conditional on the models.

Being a discrete probability distribution on consequences, lottery $m_\theta \circ f^{-1}$ can be represented as

$$m_\theta \circ f^{-1} \equiv \left\{ \left\langle m_\theta \left( f^{-1}(x) \right), x \right\rangle : x \in f(S) \right\} = \left\{ \left\langle m_\theta (S_n), x_n \right\rangle : n = 1, ..., N \right\} = \left\{ \left\langle (\phi_\theta)_n, x_n \right\rangle : n = 1, ..., N \right\}.$$

**Glossary** We have obtained a full correspondence between Kaplan and Garrick's probability of frequency format and the decision theory framework of CMMM. Here is the conversion table

| KG probability of frequency format | Classical decision theory framework |
|---|---|
| $H = \{\langle S_n, x_n \rangle : n = 1, ..., N\}$ | $f(s) = x_n \quad s \in S_n \ (n = 1, ..., N)$ |
| $\phi_\theta$ | $m_\theta$ |
| $(\phi_\theta)_n$ | $m_\theta (S_n)$ |
| $\rho_\theta$ | $\mu (m_\theta)$ |
| $\rho_n (\phi_n) = \{\langle \rho_\theta, (\phi_\theta)_n \rangle : \theta \in \Theta\}$ | $\mu^{S_n} = \sum_{\theta \in \Theta} \mu (m_\theta) \delta_{m_\theta(S_n)}$ |
| $R = \langle S_n, \rho_n (\phi_n), x_n \rangle$ | $(f \mid \mu)$ |
| $\Phi_\theta$ | $G_f^\theta$ |

## 4.1 Special embeddings

As observed at the end of Section 2, to each risk in probability of frequency format $R = \{\langle S_n, \rho_n (\phi_n), x_n \rangle : n = 1, ..., N\}$ it naturally corresponds a reduced risk in probability format $\bar{R} = \{\langle S_n, \bar{p}_n, x_n \rangle : n = 1, ..., N\}$ where, for each $n = 1, ..., N$,

$$\bar{p}_n = \sum_{\theta \in \Theta} (\phi_\theta)_n \rho_\theta = \sum_{\theta \in \Theta} m_\theta (S_n) \mu (m_\theta) = P_\mu (S_n).$$

This establishes the correspondence between, Kaplan and Garrick's probability format and the decision theory framework of Savage where only the predictive probability $P_\mu$ is considered (see again the discussion in CMMM).

On the other hand, Kaplan and Garrick's frequency format corresponds to the degenerate probability of frequency format in which $\Theta$ is a singleton, so that also $\mathcal{M} = \{m\}$ is a singleton, and the frequency $\phi_n$ of each scenario $S_n$ is simply $m(S_n)$. Formally

$$\begin{aligned} R &= \{\langle S_n, \phi_n, x_n \rangle : n = 1, ..., N\} \\ &= \{\langle S_n, m(S_n), x_n \rangle : n = 1, ..., N\} \\ &= \left\{\left\langle f^{-1}(x), m\left(f^{-1}(x)\right), x \right\rangle : x \in f(S)\right\} \equiv (f \mid m) \end{aligned}$$

so that Kaplan and Garrick's frequency format corresponds to the decision theory framework of von Neumann and Morgenstern where an objective probability $m$ over states is given.

## 5 Classical Expected Utility

At this point we consider the characterizing ingredient of decision theory: preferences. In this framework, DM's preferences are represented by a binary relation $\succsim$ defined over the set of all acts. In particular, we write $f \succsim g$ if act $f$ is (weakly) preferred to act $g$. Denoting by $x$ both a consequence and the constant act that delivers $x$ in every state, we can write $x \succsim y$. The preference $\succsim$ is thus able to rank also consequences.

Among acts, bets play a special role since it is through them that subjective probabilities are elicited. In particular, consider any two consequences $x$ and $y$, and let $x \succ y$, that is, $x$ be preferred to $y$. Then, we denote by $xEy$ the bet on event $E$ that pays the best consequence $x$ if $E$ obtains and $y$ otherwise. Given any two events $E$ and $F$, a preference $xEy \succsim xFy$ reveals that the DM considers $E$ more likely than $F$. If, moreover, $\succsim$ satisfies the Savage axioms, then there exists a unique subjective probability $P$ over $\Sigma$ such that

$$xEy \succsim xFy \Longleftrightarrow P(E) \geq P(F).$$

Again, the subjective probability $P$ quantifies the DM beliefs and is elicited via betting behavior. Now when a decision problem $(\mathcal{A}, S, X, \mathcal{M})$ is considered, it seems natural to assume betting behavior to be consistent with the statistical datum $\mathcal{M}$, that is,

$$m(E) \geq m(F) \quad \forall m \in \mathcal{M} \quad \Longrightarrow \quad xEy \succsim xFy \tag{12}$$

for all events $E$ and $F$ and consequences $x \succ y$. If event $E$ is more likely than event $F$ for all possible probabilistic models, then the DM prefers to bet on $E$ than on $F$ accordingly.

As CMMM show, if $\succsim$ satisfies Savage's axioms and the consistency condition (12), then there exists a subjective prior probability $\mu$ on $\mathcal{M}$, representing epistemic uncertainty, such that acts are ranked according to the criterion

$$U(f \mid \mu) = \sum_{m \in \mathcal{M}} \left( \sum_{n=1}^{N} u(f(S_n)) \, m(S_n) \right) \mu(m) \tag{13}$$

where:

- $u : X \to \mathbb{R}$ is a von Neumann-Morgenstern utility function that captures risk attitudes (i.e., attitudes towards aleatory uncertainty);
- $U(f \mid m) = \sum_{n=1}^{N} u(f(S_n)) \, m(S_n)$ is the von Neumann-Morgenstern expected utility of $f$ conditional on model $m$.

Since positing a collection $\mathcal{M}$ of models is a central tenet of classical statistics, CMMM call *Classical Subjective Expected Utility* the representation (13). The corresponding choice rule is maximization of the average expected utility of act $f$ with respect to the probabilistic models posited by the DM according to the epistemic distribution $\mu$. Or, equivalently,

$$f \succsim g \quad \Longleftrightarrow \quad U(f \mid \mu) \geq U(g \mid \mu).$$

But now, using our glossary, we have that

$$U(f \mid \mu) = \sum_{\theta \in \Theta} \left( \sum_{n=1}^{N} u(x_n) \, (\phi_\theta)_n \right) \rho_\theta = \sum_{\theta \in \Theta} \rho_\theta U_\theta = U(R)$$

where $U_\theta = \sum_{n=1}^{N} u(x_n) \, (\phi_\theta)_n = U(f \mid m_\theta)$ is the von Neumann-Morgenstern expected utility of risk $R_\theta = \{\langle S_n, (\phi_\theta)_n, x_n \rangle : n = 1, ..., N\}$. That is, $U(f \mid \mu)$ is the decision theoretic translation of the expected utility $U(R)$ that KG (p. 23) use to evaluate risk $R$, a few lines before closing their pioneering paper. This observation concludes our translation work. In the next section we briefly sketch how our exercise can bring novel tools to risk informed decision making in the presence of both epistemic and aleatory uncertainty.

# 6 Robust criteria

The idea of including aversion to epistemic uncertainty in risk analysis dates back to Paté-Cornell and Fischbeck (1992). Their crucial intuition is that the decision theoretic counterpart of aversion to epistemic uncertainty is *ambiguity aversion*; which, in terms of statistical decision making, corresponds to a preference for robust procedures (see CMMMb). On the other hand, representation (13) presupposes that the expected utilities

$$U(f \mid m) \qquad m \in \mathcal{M}$$

conditional on probabilistic models are averaged to produce the unconditional evaluation

$$U(f \mid \mu) = \sum_{m \in \mathcal{M}} U(f \mid m)\, \mu(m) \,. \tag{14}$$

This means that the DM is neutral about "model risk," the risk involved in not knowing the true probabilistic model $m$.[9] More precisely, the DM is neutral to epistemic uncertainty, or *ambiguity neutral*. For example, (14) implies that facing an hazard $f$ with conditional expected losses

$$-U(f \mid m) \qquad m \in \mathcal{M}$$

*that depend on the true model* is indifferent to facing an hazard $g$ with conditional expected losses

$$-U(g \mid m) = -\sum_{q \in \mathcal{M}} U(f \mid q)\, \mu(q) \qquad m \in M$$

*that do not depend on the true model.* On the other hand, in many circumstances, there may be a premium in the adoption of policies that are not affected by model misspecification. Especially, when catastrophic risks are considered, it may be normatively compelling to be averse to epistemic uncertainty, or *ambiguity averse*.[10]

To distinguish attitudes toward aleatory and epistemic uncertainties, Klibanoff et al. (2005) introduce, in decision theory, preferences represented by

$$V(f \mid \mu) = v^{-1}\left(\sum_{m \in \mathcal{M}} v(U(f \mid m))\,\mu(m)\right) \tag{15}$$

where $v$ is now a strictly increasing and continuous function that represents the attitudes of the DM towards epistemic uncertainty.[11] For example a decision maker ($\mathrm{DM}_1$) is more ambiguity averse than another decision maker ($\mathrm{DM}_2$) if and only if $v_1$ is more concave than $v_2$. Klibanoff et al. (2005) call criterion (15) the *smooth ambiguity criterion*. Criterion (14) is the special case in which the DM is neutral to epistemic uncertainty, that is $v$ is affine.

But now, using our glossary, we have that

$$V(R) = v^{-1}\left(\sum_{\theta \in \Theta} v\left(\sum_{n=1}^{N} u(x_n)(\phi_\theta)_n\right)\rho_\theta\right)$$

when risk $R = \{\langle S_n, \rho_n(\phi_n), x_n\rangle : n = 1, ..., N\}$ is considered, thus yielding the risk triplet counterpart of the criterion proposed by Paté-Cornell and Fischbeck (1992, p. 208) in a (von Neumann-Morgenstern) compound-lotteries environment.

Finally, it is well known that the smooth ambiguity criterion corresponding to $v(t) = -e^{-\lambda t}$ for all $t \in \mathbb{R}$, with constant absolute ambiguity aversion coefficient $\lambda > 0$, can be written as

$$V(f \mid \mu) = -\frac{1}{\lambda}\ln\left(\sum_{m \in \mathcal{M}} e^{-\lambda U(f|m)}\mu(m)\right) = \inf_{\nu \ll \mu}\left(\sum_{m \in \mathcal{M}} U(f \mid m)\,\nu(m) + \frac{1}{\lambda}K(\nu||\mu)\right) \tag{16}$$

where $K$ is the Kullback-Leibler divergence. This suggests considering, in the spirit of Ben-Tal et al. (1991), general robust criteria of the form

$$V(f \mid \mu) = \inf_{\nu \ll \mu}\left(\sum_{m \in \mathcal{M}} U(f \mid m)\,\nu(m) + D(\nu||\mu)\right) \tag{17}$$

where $D$ is a generic divergence between priors, that is, a function

$$D : \Delta(\mathcal{M}) \times \Delta(\mathcal{M}) \to [0, \infty]$$

such that $D(\cdot||\mu)$ is convex and $D(\mu||\mu) = 0$ for every $\mu$.[12] Maccheroni et al. (2006) and CMMMb show how $D(\cdot||\cdot)$ captures ambiguity attitudes in a simple way: $\mathrm{DM}_1$ is *more ambiguity averse* than $\mathrm{DM}_2$ if $D_1(\cdot||\mu) \leq D_2(\cdot||\mu)$ for every $\mu \in \Delta(\mathcal{M})$.

We conclude by observing that, when $D$ is identically equal to 0, that is, assuming maximal ambiguity aversion,[13] we obtain

$$W(f \mid \mu) = \min_{m \in \mathcal{M}_\mu} U(f \mid m) \tag{18}$$

where $\mathcal{M}_\mu = \{m \in \mathcal{M} : \mu(m) > 0\}$. This is the fundamental robust criterion formulated by Gilboa and Schmeidler (1989).

**Example 2** *Consider a binary risk in probability of frequency format*

$$R = \{\langle S_{\text{failure}}, \{\langle \rho_\theta, (\phi_\theta)_{\text{failure}} \rangle : \theta \in \Theta\}, a - c \rangle, \langle S_{\text{safe}}, \{\langle \rho_\theta, (\phi_\theta)_{\text{safe}} \rangle : \theta \in \Theta\}, a \rangle\}$$

*that corresponds to an hazard of the type considered in Example 1. Without loss of generality assume that* $u(a-c) = 0$, $u(a) = 1$, *and* $\rho_\theta > 0$ *for all* $\theta \in \Theta$. *Then:*

– *eq. (14) prescribes*

$$U(R) = \sum_{\theta \in \Theta} \rho_\theta (\phi_\theta)_{\text{safe}}$$

*that is, it evaluates risk through the arithmetic mean-probability of success;*

– *eq. (16) prescribes*

$$V_\lambda(R) = -\frac{1}{\lambda} \ln \left( \sum_{\theta \in \Theta} \rho_\theta e^{-\lambda (\phi_\theta)_{\text{safe}}} \right)$$

*that is, it evaluates risk through the exponential mean-probability of success;*

– *eq. (18) prescribes*

$$W(R) = \min_{\theta \in \Theta} (\phi_\theta)_{\text{safe}}$$

*that is, it evaluates risk through the minimal probability of success.*

*In a nutshell,* $U(R)$ *corresponds to Bayesian risk assessment while* $W(R)$ *embodies the precautionary principle. But more is true:* $V_\lambda(R) \in (W(R), U(R))$ *for all* $\lambda > 0$ *and it monotonically* *decreases in $\lambda$, while, passing to the limit,*

$$W(R) = \lim_{\lambda \to \infty} V_\lambda(R) \qquad \textit{and} \qquad U(R) = \lim_{\lambda \to 0} V_\lambda(R).$$

*Thus, if we define, as usual in risk analysis,* acceptable *a risk for which the probability of success is as high as reasonably practicable, say $\alpha$, and $\alpha \in (W(R), U(R))$, then acceptability will depend on the value of $\lambda$. Specifically, $R$ will be acceptable if and only if*

$$-\frac{1}{\lambda} \ln \left( \sum_{\theta \in \Theta} \rho_\theta e^{-\lambda (\phi_\theta)_{\text{safe}}} \right) \geq \alpha.$$

*This inequality implicitly determines an ambiguity aversion coefficient $\lambda_\alpha \in (0, \infty)$ such that $R$ is deemed acceptable for all $\lambda \leq \lambda_\alpha$ and unacceptable otherwise. In words, acceptability will depend on the degree of aversion to epistemic uncertainty of the decision maker.*

This overview of criteria shows how robust decision making yields a reconciliation of risk analysis and the precautionary principle, which together form the basis of present day risk governance (Starr, 2003). This seems especially important in the presence of epistemic uncertainty, the relevance of which has been appreciated – in the last two decades – in all areas of risk analysis.[14]

## 7 The use of decision making criteria in risk analysis

There are multiple ways in which the criteria we discussed above can be used in risk analysis. In the approach which is typical of economic analysis, a unique criterion $V_1$ is used to choose among alternative acts. The act which maximizes $V_1$ is chosen among all the available ones. On the other hand, in risk analysis, it may also be natural to think of a multiple-stages and/or a multiple-criteria approach.

In a multiple-stage approach, first, a set of acceptable risks is determined by setting one (or more) thresholds. Say, an available act is *acceptable* if and only if it satisfies the essential requirements $V_1(f) \geq \alpha_1$, $V_2(f) \geq \alpha_2$, ... , $V_k(f) \geq \alpha_k$. Second, an act is chosen, among those acceptable, by maximizing a possibly different criterion, say $V_{k+1}$.

In a multiple-criteria approach (Figueira et al., 2005), alternatives may be ranked on the basis of several criteria at the same time. Sometimes, it could be desirable not to evaluate alternatives according to a synthetic number, obtained by aggregating the possibly different features of an act by means of a single function. In those cases, it is possible to associate a specific ranking criterion to every such feature and leaving the aggregation of them to a later stage. This framing of the problem allows, on top of other things, to disentangle objective and subjective evaluations. For example, it may be the case that a team of experts provides a ranking of alternatives based on a number of technical characteristics, leaving the final decision maker (may it be a single individual or a committee) elaborate the resultances according to a further criterion.

In particular, it follows that multiple-criteria approaches can be used also to deal with situations in which the utility of the ultimate decision maker is unknown or non-unique, as it may be the case in group decision making, where the decision protocol widely varies in the practice.

A detailed analysis of the multiple-criteria approaches to risk analysis is the object of current research, but beyond the scope of this work.

# 8 Probabilistic Risk Assessment

One of the most general (and challenging) risk analysis exercises is represented by Probabilistic Risk Assessment for a nuclear plant (PRA). Here we follow the guidelines of the U.S. NRC, verbatim. Some parts are directly reported from their official website.[15]

- A Level 1 PRA estimates the frequency of accidents that cause damage to the nuclear reactor core.

- A Level 2 PRA, which starts with the Level 1 core damage accidents, estimates the frequency of accidents that release radioactivity from the nuclear power plant.

- A Level 3 PRA, which starts with the Level 2 radioactivity release accidents, estimates the consequences in terms of injury to the public and damage to the environment.

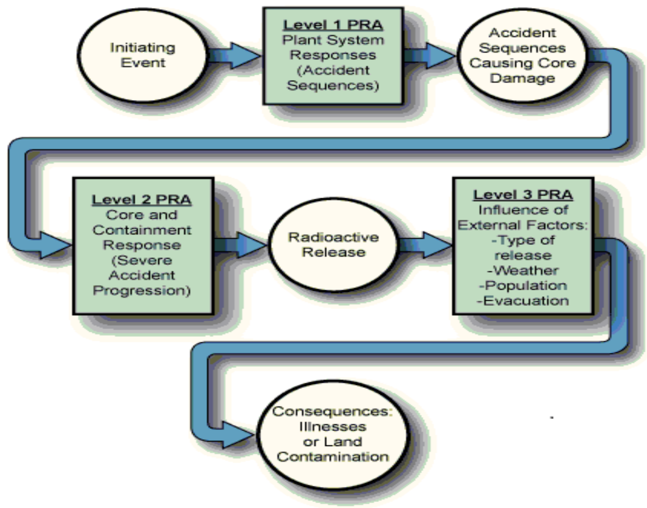


Figure 5. U.S. NRC Probabilistic Risk Assessment

It is worth to notice that the set of available alternatives depends on the application at hand. For instance, we might be evaluating competing designs for a new plant, or changes to design-bases, or operation protocols, and so on and so forth. Each of these decisions impacts how the plant sustains the structural challenge of an external (initiating) event in alternative ways, and will be briefly called design $d$.

Furthermore, as from Der Kiureghian and Ditlevsen (2009), the proper specification of the state spaces is a choice of the modeler and hence entails some degree of arbitrariness. However, the U.S. NRC and other regulatory authorities have provided guidance of what are, at each level of the analysis, the elements of interest.

**Level 1 PRA** A Level 1 PRA models the various plant responses to an event that challenges plant operation. The plant response paths are called *accident sequences*. A challenge to plant operation is called an *initiating event*. Some pairs (initiating event, accident sequence) will result in a safe recovery and some will result in reactor core damage. Thus, the state space consists of pairs

$$s_1 = (\text{initiating event, accident sequence}) \in I \times A$$

and, as a mnemonic, we set $S_1 = I \times A$. Examples of initiating events are earthquakes, terrorist attacks, and snowfalls. Here, the term initiating event is unfortunate: in a decision theoretic framework, the same object should be called "initiating state." A possible reconciliation of the two terminologies is obtained by considering singleton initiating events, which effectively correspond to initiating states. Moreover, typically, in PRA the initiating event and the design of the plant are fixed. That is, $I = \{i\}$ and only one design $d$ is considered. Since we want to compare different designs in terms of their ability to withstand possible initiating events, we allow for a nonsingleton $I$.

Although, in general, one can allow for different degrees of core damage, a binary assignment, core damage or safe recovery, is used in the practice. Therefore, the acts/hazards considered at this level, and depending on the design $d$ of the plant, are of the form

$$f'_d(s_1) = \begin{cases} \text{core damage} & s_1 \in S^d_{\text{severe}} \\ \\ \text{safe recovery} & s_1 \in S^d_{\text{safe}} \end{cases}$$

In this case, the act is uniquely determined by the event $S^d_{\text{severe}}$ which is the set of all *severe accidents* (see the first line of Figure 5).

The Level 1 PRA also provides the core damage frequencies (or, more in general, likelihoods) thus allowing the transition from acts/hazards of the form

$$f'_d \equiv \left\{ \left\langle S^d_{\text{severe}}, \text{core damage} \right\rangle, \left\langle S^d_{\text{safe}}, \text{safe recovery} \right\rangle \right\} = H_d$$

to risks of the form

$$R_d = \left\{ \left\langle S^d_{\text{severe}}, \left\{ \left\langle \rho_\theta, (\phi_\theta)^d_{\text{severe}} \right\rangle : \theta \in \Theta \right\}, \text{core damage} \right\rangle, \left\langle S^d_{\text{safe}}, \left\{ \left\langle \rho_\theta, (\phi_\theta)^d_{\text{safe}} \right\rangle : \theta \in \Theta \right\}, \text{safe recovery} \right\rangle \right\}.$$

Note that:

- core damage frequencies correspond to frequencies of the event $S^d_{\text{severe}}$;
- the risks obtained at this level are mathematically identical to those discussed in Example 2, hence the analysis we carried on there directly applies to this level of risk analysis.

**Level 2 PRA** A Level 2 PRA models the plant's response to the Level 1 PRA severe accidents, and analyzes the progression of an accident by considering how the containment structures and systems respond to the accident, which varies based on the initial status of the structure or system and its ability to withstand the harsh accident environment. Once the containment response is characterized, the analyst can determine the amount and type of radioactivity released from the containment. Thus, the Level 2 PRA state space *conditional on design* $d$, consists of pairs

$$s_2^d = (\text{severe accident, containment response}) \in S_{\text{severe}}^d \times C.$$

and consequences are radioactivity releases $Q$. Hence, conditional on $S_{\text{severe}}^d$, an act considered at this level is a map

$$f_d'' : S_{\text{severe}}^d \times C \to Q.$$

Since we want to compare different designs $d$, which induce different partitions $\left\{S_{\text{severe}}^d, S_{\text{safe}}^d\right\}$ of $I \times A$, we complete the act by assigning 0 release to non *severe accidents*, so that

$$f_d'' : S_1 \times C \to Q$$

with the convention $f_d''(s_2) = 0$ if $s_2 \in S_{\text{safe}}^d \times C$, thus $S_2 = S_1 \times C$ (see again Figure 5).

Level 2 PRA also provides the release frequencies (or, more in general, likelihoods) by analyzing the frequencies of the relevant events in $S_2$ (see the discussion of Level 1 PRA).

**Level 3 PRA** A Level 3 PRA considers the effects of the radioactive material released in a severe accident (which, by Level 2 PRA, depends on the realized state $s_2$, and the chosen design $d$).

Specifically, this level is concerned with *health effects* (such as short-term injuries or long-term cancers) resulting from the radiation doses to the population around the plant and *land contamination* resulting from radioactive material released in the accident. These consequences are estimated based on the characteristics of the radioactivity release calculated by the Level 2 PRA. They also depend on several other factors such as the *population in the plant vicinity*, the *evacuation conditions* and the path of the radioactive plume, which, in turn, is affected by *weather* *conditions.* Thus, the Level 3 PRA state space, consists of tuples

$$s_3 = (s_2, \text{weather}, \text{population}, \text{evacuation}) \in S_2 \times W \times Z \times E$$

briefly $S_2 \times W \times Z \times E = S_3$. Note that $S_3$ is the largest state space because all the unknown factors jointly determine the final (relevant) consequences. Recall

| | | |
|---|---|---|
| PRA1 | $S_1 = I \times A$ | only severe accidents determine core damage, |
| PRA2 | $S_2 = S_1 \times C$ | also containment response is relevant for radioactive releases, |
| PRA3 | $S_3 = S_2 \times W \times Z \times E$ | additional environmental conditions lead to final consequences. |

Denoting by $J \times L$ the pairs $(j, l)$ of health effects and land contamination levels, we are now ready to write the level 3 acts:

$$f_d^{'''} : S_3 \to J \times L$$

that associate complete states to relevant consequences. Note that here consequences are multiattribute vectors.[16]

Finally, like in the previous two cases, Level 3 PRA estimates the likelihoods of events in $S_3$, which ultimately determine the likelihoods of the consequences that are relevant for the DM's decision. The construction we just described leads us to a complete specification of risk triplets (as detailed in Section 4 and exemplified for Level 1 PRA), which are precisely the object of the analysis performed by the U.S. NRC (What can go wrong, how likely is it, and what are the consequences?), *sic.*

The added value of the above exercise is framing the PRA in a setup that allows to compare different designs, for example, by evaluating the obtained third level acts through some of the criteria discussed in the previous sections.

## 9 Conclusions

As promised, a formal bridge between risk analysis and decision theory has been built and tested on the benchmark of nuclear PRA. The hope is that it will foster the cross-fertilization of the two fields. Indeed, given the depht and scope of the two underlying subjects, the possibilities of mutual enrichment are substantial.

An immediate benefit of our analysis is providing a structured decision making framework to the regulator. Indeed, *whether a decision issue is straightforward or difficult, a nuclear regulator will benefit by having a structured decision-making framework* (Murley, 2005). Another connected advantage is that it makes starkly explicit the assumptions underlying decision making, thus allowing more careful considerations upon them (Apostolakis, 2014).

## Notes

[1] For an early and critical review about risks and benefits of technological systems, we refer to Starr (1969).

[2] The problem was already clear at the time of KG. As they state: *one often hears people say that we cannot use probability because we have insufficient data, in light of our current definitions, we see that this is a misunderstanding. When one has insufficient data, there is nothing else one can do but use probability.*

[3] In particular, deterministic models correspond to degenerate probabilistic ones.

[4] See Gilboa and Marinacci (2013) for a recent and comprehensive review.

[5] Though here $\bar{p}_n$ is a reduced probability that arises in a repetitive situation.

[6] In what follows, events are always elements of $\Sigma$; for example, Borel subsets of the real line.

[7] There is a little abuse of notation here, justified by the fact that for every $s_n, t_n \in S_n$ we have $f(s_n) = f(t_n)$, so that we can write $f(S_n)$ instead of $f(s_n)$.

[8] In decision theory, lottery is a synonym of discrete probability distribution over consequences.

[9] Paté-Cornell and Fischbeck (1992, p. 208) call $m$ a *model of the world.*

[10] Paté-Cornell and Fischbeck (1992, p. 203) refer to *severe core damage of a nuclear reactor* and *submit that rationality can and should be extended to include willingness to pay a systematic premium to replace ... epistemic uncertainties ... by aleatory uncertaint[ies].*

[11] While $V$ represents $\succsim$ in the sense that $f \succsim g \iff V(f \mid \mu) \geq V(g \mid \mu)$.

[12] We denote by $\Delta(\mathcal{M})$ the simplex of all priors on $\mathcal{M}$.

[13] For example, letting the ambiguity aversion coefficient $\lambda$ go to $\infty$ in (16).

[14] Some applications include Bedford (2013), Helton and Johnson (2011), Hora (1996), Marzocchi et al. (2004) and Mert and Thieken (2005).

[15] http://www.nrc.gov/about-nrc/regulatory/risk-informed/pra.html

[16] According to the U.S. NRC description of this level of analysis one can assume that the consequence vector $f_d'''(s_3)$ depends from the $s_2$ component of $s_3 = (s_2, w, z, e)$ only through $f_d''(s_2)$ which describes the corresponding radioactivity release.

## Acknowledgments

We thank Elisabeth Paté-Cornell, Salvatore Greco, and Curtis Smith for very insightful discussions, and the audiences at the DAS Conference (June 16-18, 2014, Washington DC) and the PSAM Conference (June 22-27, 2014, Honolulu) for useful comments.

The financial support of MIUR (PRIN grant 20103S5RN3_005) and of the AXA Research Fund is gratefully acknowledged.

## References

Althaus, C., A disciplinary perspective on the epistemological status of risk, *Risk Analysis*, **25**, 567–88, 2005.

Apostolakis, G. E., The concept of probability in safety assessments of technological systems, *Science*, **250**, 1359–1364, 1990.

Apostolakis, G. E., How useful is quantitative risk assessment?, *Risk Analysis*, **24**, 515–520, 2004.

Apostolakis, G. E., Beware of the assumptions: decision making and statistics, PSAM12, plenary lecture, 2014.

Apostolakis, G. E. and Pickett, S., Deliberation: Integrating analytical results into environmental decisions involving multiple stakeholders, *Risk Analysis*, **18**, 621–634, 1998.

Aven, T., Foundational issues in risk assessment and risk management, *Risk Analysis*, **32**, 1647–1656, 2012.

Aven, T., The risk concept - historical and recent development trends, *Reliability Engineering & System Safety*, **99**, 33–44, 2012b.

Borgonovo, E. and Smith, C., A study of interactions in the risk assessment of complex engineering systems: An application to space psa, *Operations Research*, **59**, 1461–1476, 2011.

Bedford, T., Decision making for group risk reduction: dealing with epistemic uncertainty, *Risk Analysis*, **33**, 1884–1898, 2013.

Ben-Tal, A., Ben-Israel, A. and Teboulle, M., Certainty equivalents and information measures: duality and extremal principles, *Journal of Mathematical Analysis and Applications*, **157**, 211–236, 1991.

Boykin, R. F., Freeman, R. and Levary, R. R., Risk assessment in a chemical storage facility, *Management Science*, **30**, 512–517, 1984.

Caruso, M., Cheok, M., Cunningham, M., Holahan, G. M., King, T., Parry, G., Ramey-Smith, A. M., Rubin, M. and Thadani, A., An approach for using risk assessment in risk-informed decisions on plant-specific changes to the licensing basis, *Reliability Engineering & System Safety*, **63**, 231–242, 1999.

Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M. and Montrucchio, L., Classical subjective expected utility, *PNAS*, **110**, 6754–6759, 2013.

Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M. and Montrucchio, L., Ambiguity and robust statistics, *Journal of Economic Theory*, **148**, 974–1049, 2013b.

Clemen, R. T. and Reilly, T., Correlations and copulas for decision and risk analysis, *Management Science*, **45**, 208–224, 1999.

Der Kiureghian, A. and Ditlevsen, O., Aleatory or epistemic? does it matter?, *Structural Safety*, **31**, 105–112, 2009.

Dillon, R., Paté-Cornell, M. E. and Guikema, S., Programmatic risk analysis for critical engineering systems under tight resouece constraints, *Operations Research*, **51**, 354–367, 2003.

Figueira, J., Greco, S. and Ehrgott, M., (eds.) *Multiple Criteria Decision Analysis: State of the Art Surveys*, Springer, 2005.

Garrick, B. J. and Kaplan, S., A decision high-level theory perspective on the disposal of radioactive waste, *Risk Analysis*, **19**, 903–913, 1999.

Gilboa, I. and Marinacci, M., Ambiguity and the Bayesian paradigm. In *Advances in Economics and Econometrics: Theory and Applications* (eds. Acemoglu, D., Arellano, M. and Dekel, E.), 179–242, Cambridge: Cambridge University Press, 2013.

Gilboa, I. and Schmeidler, D., Maxmin expected utility with a non-unique prior. *Journal of mathematical economics*, **18**, 141–153, 1989.

Helton, J. and Johnson, J., Quantification of margins and uncertainties: alternative representations of epistemic uncertainty, *Reliability Engineering & System Safety*, **96**, 1034–1052, 2011.

Howard, R., Decision analysis: Practice and promise, *Management Science*, **34**, 679–695, 1988.

Hora, S. C., Aleatory and epistemic uncertainty in probability elicitation with an example from hazardous waste management, *Reliability Engineering & System Safety*, **54**, 217–223, 1996.

Kaplan, S., The words of risk analysis, *Risk Analysis*, **17**, 407–417, 1997.

Kaplan, S. and Garrick, B. J., On the quantitative definition of risk, *Risk Analysis* **1**, 1–28, 1981.

Klibanoff, P., Marinacci, M. and Mukerji, S., A smooth model of decision making under ambiguity, *Econometrica*, 73, 1849–1892, 2005.

Maccheroni, F., Marinacci, M. and Rustichini, A., Ambiguity aversion, robustness, and the variational representation of preferences, *Econometrica*, 74, 1447–1498, 2006.

Marzocchi, W., Sandri, L., Gasparini, P., Newhall, C. and Boschi, E., Quantifying probabilities of volcanic events: the example of volcanic hazard at Mount Vesuvius, *Journal of Geophysical Research: Solid Earth*, **109**, 1978–2012, 2004.

Marzocchi, W., Newhall, C. and Woo, G., The scientific management of volcanic crises, *Journal of Volcanology and Geothermal Research*, **247-248**, 181–189, 2012.

Murley, T., Nuclear regulatory decision making, *OECD report*, 2005.

Merz, B. and Thieken, A. H., Separating natural and epistemic uncertainty in flood frequency analysis, *Journal of Hydrology*, **309**, 114–132, 2005.

Paté-Cornell, M. E., Conditional uncertainty analysis and implications for decision making: The case of WIPP, *Risk Analysis*, **19**, 995–1002, 1999.

Paté-Cornell, M. E. and Dillon, R., The respective roles of risk and decision analyses in decision support, *Decision Analysis*, **3**, 220–232, 2006.

Paté-Cornell, M. E. and Fischbeck, P. S., Aversion to epistemic uncertainties in rational decision making: effects on engineering risk management. In *Risk-Based Decision Making in Water Resources V* (eds. Y. Haimes and D.A. Moser), 200–218, New York: ASCE, 1992.

Pratt, J., Raiffa, H. and Schlaifer, R., *Statistical Decision Theory*, third ed., Cambridge: MIT Press, 1995.

Savage, L., *The foundations of Statistics*, New York: Wiley and Sons, 1954.

Schmeidler, D., Subjective probability and expected utility without additivity, *Econometrica*, **57**, 571–587, 1989.

Smith, J. and von Winterfeldt, D., Decision analysis in management science, *Management Science*, **50**, 561–574, 2004.

Stamatelatos, M., Dezfuli, H. and Others, Nasa probabilistic risk assessment procedures guide for nasa managers and practitioners, *NASA/SP-2011-3421,*third ed., 2011.

Starr, C., Social benefit versus technological risk, *Science*, **165**, 1232–1238, 1969.

Starr, C., The precautionary principle versus risk analysis, *Risk Analysis*, **23**, 1–3, 2003.

von Neumann, J. and Morgenstern, O., *Theory of Games and Economic Behavior*, second ed., Princeton: Princeton University Press, 1947.

Wald, A., *Statistical decision functions*, New York: John Wiley and Sons, 1950.

Winkler, R., Uncertainty in probabilistic risk assessment, *Reliability Engineering & System Safety*, **54**,127–132, 1996.