institut interdisciplinaire de l'innovation
WORKING PAPERS
Documents de travail i3

Ambiguity aversion and the expected cost of rare energy disasters

AN APPLICATION TO NUCLEAR POWER ACCIDENTS

Romain Bizet
MINES ParisTech, PSL Research University, CERNA - Centre for industrial economics, i3, CNRS UMR 9217
romain.bizet@mines-paristech.fr

François Lévêque
MINES ParisTech, PSL Research University, CERNA - Centre for industrial economics, i3, CNRS UMR 9217
francois.leveque@mines-paristech.fr

Working Paper 16-CER-01
Avril, 2016

Pour citer ce papier / How to cite this paper : Bizet, R. & Lévêque, F. (2016) Ambiguity aversion and the expected cost of rare energy disasters: An application to nuclear power accidents. i3 Working Papers Series, 16-CER-01.

L’institut interdisciplinaire de l’innovation

(UMR 9217) a été créé en 2012. Il rassemble :

- les équipes de recherche de MINES ParisTech en économie (CERNA), gestion (CGS) et sociologie (CSI),
- celles du Département Sciences Economiques et Sociales (DSES) de Télécom ParisTech,
- ainsi que le Centre de recherche en gestion (CRG) de l’École polytechnique,

soit plus de 200 personnes dont une soixantaine d’enseignants chercheurs permanents.

L’institut développe une recherche de haut niveau conciliant excellence académique et pertinence pour les utilisateurs de recherche. Par ses activités de recherche et de formation, i3 participe à relever les grands défis de l’heure : la diffusion des technologies de l’information, la santé, l’innovation, l’énergie et le développement durable. Ces activités s’organisent autour de quatre axes :

- Transformations de l’entreprise innovante
- Théories et modèles de la conception
- Régulations de l’innovation
- Usages, participation et démocratisation de l’innovation

Pour plus d’information : http://www.i-3.fr/

The Interdisciplinary Institute of Innovation

(UMR 9217) was founded in 2012. It brings together:

- the MINES ParisTech economics, management and sociology research teams (from the CERNA, CGS and CSI),
- those of the Department of Economics and Social Science (DSES) at Télécom ParisTech,
- and the Management Research Center (CRG) at Ecole Polytechnique,

meaning more than 200 people, including 60 permanent academic researchers.

i3 develops a high-level research, combining academic excellence and relevance for the end users of research. Through its teaching and research activities, i3 takes an active part in addressing the main current challenges: the diffusion of communication technologies, health, innovation, energy and sustainable development. These activities are organized around four main topics:

- Transformations of innovating firms
- Theories and models of design
- Regulations of innovation
- Uses, participation and democratization of innovation

For more information: http://www.i-3.fr/

Ambiguity Aversion and the expected cost of Rare Energy Disasters: An Application to Nuclear Power Accidents.*

Romain Bizet[†1] and François Lévêque[1]

[1]Mines ParisTech, PSL - Research University, CERNA - Centre for Industrial Economics, i3, CNRS UMR 9217

April 11, 2016

Abstract

Assessing the risks of rare disasters due to the production of energy is paramount when making energy policy decisions. Yet, the costs associated with these risks are most often not calculable due to the high uncertainties that characterize their potential consequences. In this paper, we propose a non-Bayesian method for the calculation of the expected cost of rare energy disasters, that accounts for the ambiguity that characterizes the probabilities of these events. Ambiguity is defined as the existence of multiple and conflicting sources of information regarding the probabilities associated with these events. We then apply this method to the particular case of nuclear accidents in new builds. Our results suggests that the upper-bound of the expected cost of such accidents is 1.7€/MWh, which is consistent with most of the recent estimates. This expected cost may rise to 7 €/MWh when the macroeconomic shock caused by a nuclear accident is taken into account.

Keywords: nuclear accidents, rare disasters, expected costs, uncertainty, ambiguity.

JEL Classification numbers: D62, D81, Q48.

1 Introduction

Not all consequences associated with energy production technologies can be considered as regular externalities. Namely, the risks associated with conventional means of production of energy are most

*We are particularly indebted to Jean-Marc Tallon for his numerous comments, suggestions and references. We would also like to thank Margaret Kyle, Pierre Fleckinger, and all the participants of the doctoral seminar of the department of economics in Mines ParisTech for their very helpful comments.

†Contact details: Mines ParisTech, 60 Boulevard Saint-Michel, 75006, Paris. romain.bizet@mines-paristech.fr

often characterized by their catastrophic consequences: nuclear accidents, oil spills or dam failures are good examples of such disasters. Moreover, energy production technologies have numerous consequences which cannot be described by a probability distribution over monetary outcomes. What are, for instance, the probabilities of dam failures or major nuclear accidents? Yet, governments have to choose which technology to use, banks have to decide which project to finance, and utilities have to determine when to shut down old plants and what to replace them by. An important underlying question of the energy debate is how should we compare the costs, the risks and the benefits associated with each energy production scenario? In other words, on which basis should choices among these alternatives be made?

Historically, cost-benefit analysis (CBA) has been used to assess the desirability of such projects. Hanley and Spash (1993), for instance, describe the rise of the use of CBA in the nineteenth century in the United States, from the evaluation of irrigation and flood-control policies, to the general assessment of projects affecting public goods, such as human health or the environment. Yet, several limits of CBA were exposed in the late 1990 by the climate change literature, such as its failures to deal rigorously with the choice of a discount rate or the fair allocation of resources, or to account properly for the probabilities associated with the outcomes of climate change and the monetary valuation of climate damage (Azar and Lindgren, 2003). More recently, Weitzman (2011) criticized the use of CBA when uncertainty[1] prevails, and argued in favour of the use of fat-tailed distributions to account for the possibility of extreme climatic disasters.

Cost-benefit analysis is based on three assumptions. It first stands on the hypothesis that among competing projects, the most desirable is the one that yields the highest expected benefit (or equivalently the lowest expected cost). Second, for these expected benefits to be calculated, CBA requires that the potential consequences of each project be known, and associated with well-defined monetary values. Finally, the existence of a probability distribution over these potential consequences is assumed. When assessed projects are characterized by rare occurrences of catastrophic consequences, the last hypothesis does not hold. In particular, rare disasters fail to satisfy the definition of objective probabilities proposed by Savage (1954): their probabilities cannot be convincingly identified with their observed frequencies of occurrence, as the repetitions of these events are neither independently nor identically distributed. Moreover, other assessments of the probabilities of rare disasters may also be available, such as public opinion or probabilistic safety analyses performed by energy producers.

[1] In this paper, following the terminology defined by Knight (1921), risk will refer to situations that can be represented as lotteries associated with known probabilities and outcomes. Uncertainty will refer to situations in which either probabilities or outcomes are vague or even unknown. Another equivalent terminology opposes *aleatory risks* (e.g. risks) with *epistemic risks* (e.g. uncertainty). See Paté-Cornell (1996) for further references.

When these sources are in sharp contradiction, which one(s) should be used to guide the choices among available energy production technologies? Calculating an expected cost based on either of these conflicting probabilities may seem like an *ad hoc* choice, rather than a rigorous method on which sound decisions may be made. This paper intends to generalize cost-benefit analysis to the study of events which cannot properly be described by a single probability distribution over monetary outcomes.

To do so, we refer to the theoretical literature dedicated to decision-making under *ambiguity*. More precisely, Ellsberg (1961) first showed that ambiguity, or the absence of knowledge regarding the probabilities of some events, had an effect on individual behaviours that was different from the influence of classical risk. As a consequence, scholars proposed various decision criteria that explicitly account for ambiguity, and for ambiguity-aversion, e.g. the attitude of people towards ambiguity. Examples of such criteria can be found in Schmeidler (1989); Gilboa and Schmeidler (1989); Bewley (2003) or Klibanoff et al. (2005). Applications of this theoretical literature were mostly developed in the finance literature, for example in Dow and Werlang (1992); Epstein and Wang (1994); Chateauneuf et al. (1996) or Epstein and Schneider (2008). Yet, ambiguity has also been applied to the study of rare disasters. Paté-Cornell (1996) suggested the use of decision processes that acknowledge ambiguity and ambiguity aversion in the study of uncertain risks; and Crès et al. (2011) proposed a way to aggregate conflicting opinions of experts regarding the consequences of climate change. Our paper describes a new method for the assessment of the expected costs of rare disasters related to the production of energy, which coincides with traditional cost-benefit analysis when objective probabilities are available, but also allows to account for multiple probability distributions when facing ambiguity. This method is based on the α-maxmin criterion proposed by Hurwicz (1951) and developed by Ghirardato et al. (2004).

Finally, we propose an application of our method to the assessment of the expected cost of major nuclear power accidents. Our contribution to this literature is twofold. By introducing ambiguity and ambiguity-aversion in the analysis of the risks associated with the use of nuclear power, this paper furthers the efforts of Eeckhoudt et al. (2000), who tried to include risk-aversion in their analysis. Moreover, we argue that our non-Bayesian[2] method is more adapted to the analysis of rare events than the statistical methods which have been applied until now to the estimation of the costs and probabilities of rare disasters. Examples of such methods can be found in Hofert and Wüthrich (2011),

[2]A definition of Bayesian decision-making can be found in Gilboa (2004). Its main three characteristics are that the probabilities associated with any state of the world are known, at least subjectively; that decision-makers use Bayes rule when they can; and that decision-makers make their decisions according to a decision rule that consist in maximizing an expected utility with respect to known probabilities. The method presented here violates the first and third characteristics.

D'Haeseleer (2013) or Escobar Rangel and Lévêque (2014)[3].

This paper is organised as follows. Section 2 will define our method for the calculation of the expected cost of rare disasters. Section 3 will present our application of this method to nuclear accidents. Section 4 will discuss some limits of our model. Section 5 will conclude and discuss some policy implications.

2 Decision-making and rare disasters

2.1 An example: the ambiguity associated with nuclear disasters

2.1.1 Past events

When considering the risks related to the use of nuclear power, the first source of knowledge regarding safety is experience: past events reveal information regarding the probability of nuclear accidents. The *objectivistic* view of probability described by Savage (1954) suggests that the frequency of an event is to be identified with its probability of occurrence if the repetition of the event is in close agreement with the mathematical concept of repeated, independently and identically distributed events. Since the beginning of the nuclear industry, several accidents have occurred over the world. Yet, this observed frequency of nuclear accidents cannot be identified with their probability of occurrence. Indeed, nuclear accidents are neither repeated, nor independently or identically distributed events.

First of all, the concept of repeated events is hard to apply to nuclear accidents. First, accident are always specific, because of what triggers them, but also because of how they impact society. Today, approximately four hundred reactors are operated throughout the world, for a total of 14.500 reactor.years of experience (IAEA, 2006). During this period, three severe accidents occurred (Three-Mile Island, Chernobyl and Fukushima), twelve core meltdowns were witnessed, and 41 events were ranked strictly higher than level 2 on the IAEA damage scale (Cochran and McKinzie, 2011). Which definition should then be used as the "repeated event"?

More importantly, nuclear accidents are not identically distributed. First, they are not identically distributed among reactors, since power stations are located at different locations. For instance, some reactors have to face a high seismic risk, while it is not the case for many others. Moreover, the design basis of reactors varies from one country to another, so do safety regulators, or the operators' levels of compliance with safety standards. These parameters entail changes in the likelihood of accident.

[3]An important caveat of our study is that we forego the question of the assessment of the damage associated with nuclear power accidents, or the analysis of the uncertainty that stems from this assessment. Our analysis is thus based on the latest existing studies that tackle this issue, a review of which is provided in the appendices.

Finally, nuclear accidents are not independently distributed. The first argument is that when an accident occurs, new safety regulations are set throughout the world, and the probability to witness another accident changes. A second argument to justify dependence is common cause: in Fukushima-Daiichi for instance, three core-melts were witnessed, but they were all caused by the same tsunami. Likewise, since some reactors share the same design basis, accidents caused by design failures cannot be considered as independent.

As accidents are rare, and neither independent nor identically distributed, their frequency cannot be identified with an objective probability. However, this identification can yet be found in some papers of the existing literature. For example, Rabl and Rabl (2013) calculate the expected cost of nuclear accidents, considering a future probability of occurrence based on the observed frequency of one major accident leading to widespread releases of radioactive materials every 25 years for a four hundred reactor fleet, e.g. a probability of 10^{-4} accidents per reactor.year. This assumption is often justified as a "precautionary" assumption: even though it is fairly sure that the identification is spurious, it is still done as it certainly consists in an overestimation of the probabilities of accidents.

The objectivistic foundation of statistics is not the only one that may allow to derive information regarding future accidents from the observation of the past. Therefore, other attempts have been made to apply modern statistical methods to the available data regarding nuclear accidents. As we mentioned in the literature review, D'Haeseleer (2013) and Escobar Rangel and Lévêque (2014) used two Bayesian revision frameworks in order to derive the probability of witnessing the next "Fukushima-like" accident, and Hofert and Wüthrich (2011) tried to fit a fat-tailed distribution on past accident data to derive the expected yearly damage due to nuclear accidents, and to estimate the optimal insurance mechanisms for the protection against nuclear disasters. In these studies, the order of magnitude for the probability of the next large-release accident ranges from 10^{-4} to 10^{-5} per reactor.year[4].

2.1.2 Probabilistic safety assessments

The second source of information regarding nuclear safety is a discipline that has been developed by the nuclear industry over the last forty years, called probabilistic safety assessment (PSA). This technique is based on numerical simulations and event-trees representing the potential causes and consequences of failures in the different systems that constitute a power station. The first PSA is known as the WASH-1400 report, and was carried out in the United States by Rasmussen (1975). The

[4]The *reactor.year* is the classical unit in which is expressed the probabilities of nuclear accident. It expresses the risk of witnessing a nuclear accident in a particular plant during a calendar year. To obtain the overall probability of accident for a whole fleet and for a given lifespan, one should multiply the probability expressed per reactor.year by the number of power plant in the fleet and by their associated lifespan.

interest of these safety assessments is twofold. On the one hand, they try to identify the sequences of events which could lead to nuclear accidents in order to improve the design of nuclear reactors. On the other hand, they can be used to describe the sequences of events following nuclear accidents, in order to calculate its potential cost, or to plan the emergency mitigation plans.

According to the IAEA (2010), three levels of PSAs are worth being distinguished. Level-1 PSAs analyse the design and operation of nuclear power plants. Sequences of events are identified according to the risk of core damage they entail. The primary objective of these assessments is to identify particular weaknesses in the reactors' systems, and to identify the chains of events which may lead to core damage. Though it never was their initial purpose, level-1 PSAs have also been used to derive expected frequencies of core meltdowns in nuclear power plants.

The second level of probabilistic safety assessments studies the evolution of core damage, focusing on the progression of fuel damage. The eventuality of radioactive releases to the environment is assessed quantitatively. These assessments have been used to calculate expected frequencies of large radioactive releases, and to provide insights into the cost-benefit analysis of large accident prevention. Here as well, these assessments are often used to derive probabilities of occurrence of accidents leading to large releases of radioactive materials. An example of such a use of PSA is provided on figure 1.

Figure 1: The results of AREVA's PSA for the Hinkley Point EPR (Source: Report from the U.K. Health and Safety Executive Nuclear)

Item	EDF and AREVA Target (per yr)	Result (per yr)
Core Damage Frequency (CDF) internal events	1×10^{-6}	2.77×10^{-7}
CDF ext hazards	5×10^{-6}	5.07×10^{-8}
CDF internal hazards	1×10^{-6}	3.98×10^{-8}
Offsite dose 0.1-1mSv	1×10^{-2}	1.4×10^{-3}
Offsite dose 1-10mSv	1×10^{-3}	1.3×10^{-5}
Offsite dose 10-100mSv	1×10^{-4}	8.8×10^{-7}
Offsite dose 100-1000mSv	1×10^{-5}	6.5×10^{-7}
Offsite dose >1000mSv	1×10^{-6}	5.6×10^{-8}
>100 Fatalities	1×10^{-7}	6×10^{-8}

The third level of probabilistic assessments tackles the issue of public safety during the aftermath of a major accident. It models the diffusion of radioactive materials in the environment. The impact of the radioactive fallout on public health or land use is estimated. Whereas level 1 and 2 PSAs are mostly used to identify safety weaknesses, level 3 PSAs are used to assess the cost and consequences of nuclear accidents, as well as the most efficient mitigation measures.

Probabilistic safety assessments thus aggregate the knowledge of operators and nuclear engineers. Yet, these studies have been criticized for being prone to “unknown unknowns” (Lévêque, 2013). In other words, PSAs cannot foresee all sequences of events that may trigger nuclear accidents. Furthermore, these studies are based on the assumption that regulators are independent (e.g. not captured), and that safety standards are perfectly enforced. This hypothesis of perfect compliance is very strong, as the Fukushima-Daiichi accident has shown that regulators can be captured and safety standards infringed Wang et al. (2013); Lévêque (2013), and as many countries still have very obscure nuclear safety regulators, who are either strongly tied to their industry (as it was the case in Japan before Fukushima), or to their government. Therefore PSAs provide numbers that do not reflect the actual safety level of a nuclear fleet, and underestimate the probability of nuclear accidents.

2.1.3 Public perceptions

A third source of knowledge regarding the risks of nuclear accidents is the probability perceived by the public. Indeed, people have subjective opinions regarding the probabilities of nuclear accidents, which allow them, for example, to decide whether to live close to a nuclear plant, or whether to vote for a pro-nuclear candidate during an election. These opinions have to be taken into account in social decisions regarding nuclear power as neglecting them can entail additional costs, such as the public resentment of a technology. For instance, the French fast-breeder prototype reactor Super-Phénix produced electricity between 1986 and 1996, before being shut-down because of intense social protests.

Similarly to past observations and probabilistic safety assessments, these perceptions are biased. Experimental psychology works showed that small probabilities tend to be overestimated, a possible explanation being for example excessive publicity (Slovic et al., 1977, 1982). Likewise, when facing dreadful events, people are prone to the *denominator neglect* heuristic (Kahneman, 2011): the potential damage is so large that one tends to overestimate its probability of occurrence. In addition to that, Fischhoff et al. (1978) showed that people preferred to subscribe to insurance against small probable losses than against large but rare losses. Therefore, because nuclear accidents are rare and dreadful, it can be argued that making decisions based on the public’s perceptions of the nuclear risk may yield over-investment in safety or suboptimal energy mixes.

To conclude this section, these sources of knowledge are conflictual, since they do not converge towards a unique probability. Probabilistic safety assessments suggest that probabilities of accidents are lower than 10^{-6} per reactor year. On the other hand, the observed frequencies of major accidents are close to one large accident every twenty-five years, which is equivalent to a probability of 10^{-4} per reactor.year for a 400-reactors fleet. Finally, subjective probabilities of nuclear accidents are hard to measure and are probably very heterogeneous among the population. Yet, some proxies of public opinions, such as press articles, or post-Fukushima declarations by State officials, may suggest that this perceived probability of accident is even larger than than both the results of PSAs or of past observations. For instance, a few months after the Japanese disaster, Jacques Repussard, senior officer at the French nuclear safety technical body (IRSN), opposed observed frequencies of accidents with targeted probabilistic safety objectives. In addition, the French newspaper *Liberation* published an article that claimed that the probability of a nuclear accident in Europe within the next decades was superior to one[5].

This variety of sources of knowledge regarding the probabilities of nuclear accidents illustrate the notion of ambiguity described in the introduction. Furthermore, we can generalize these remarks to other rare technological disasters. Indeed, the first argument regarding past events holds as long as the occurrences of some rare disaster are rare (e.g. not repeated), affected by local idiosyncrasies (e.g. not identically distributed), and result in technological safety upgrades (e.g. not independent). The second argument holds as long as PSAs are carried out, which is the case for most conventional energy production technologies prone to rare disasters[6]. Finally, the third argument was very general: it is sufficient for an event to be rare and dreadful for its probability to be misperceived by individuals. Therefore, ambiguity may also characterize other technological disasters, such as oil spills or dam failures.

2.2 Making good decisions under ambiguity

2.2.1 The decision criterion

How then should we make decisions when such ambiguity exists? Indeed, let's assume that a public evaluator needs to choose between several energy policies, which involve ambiguous outcomes, such as nuclear accidents. These policies may be choices of an electricity production technology mix, or several levels of safety standards for energy producers, or different laws regarding environmental pollution. In any case, if the public evaluator is aware of the multiplicity of probability distributions available, and of their respective biases, he may not want to resort to the use of a cost-benefit analysis based on

[5] References of these declarations (in French) can be found here, and here

[6] For instance, in Canada, probabilistic safety analyses of oil spill risks have been performed in 2013 (WSP Canada Inc., 2013). Likewise, France and the U.S. use probabilistic safety analyses for the assessment of dam safety. Evidence for such practices can for instance be found in the guide of best practices proposed by the U.S. Department of Interior (Bureau of Reclamation, 2015).

only one of these priors. Indeed, decisions based on biased probabilities could result in, for instance, wrong levels of safety standards, sub-optimal technology mixes, or inefficient phase-out plans for ageing energy-production facilities. Therefore, a method that acknowledges this multiplicity of information is needed in order to provide better decision guidelines.

As Ellsberg (1961) seminal paradox showed that risk and ambiguity had different impacts on individual behaviours, decision-theorists have developed representations of individual preferences that are affected by ambiguity. One of these representations was proposed by Ghirardato et al. (2004) (GMM in the following), and is known as the α-maxmin rule of decision-making under ambiguity. Intuitively, according to this criterion, one can define the expected cost of a rare disaster as the weighted average of its expected cost under a worst-case probability distribution and its expected cost under a best-case probability distribution. The weight α assigned to the worst-case expected cost represents the attitude of the decision-maker towards ambiguity. Above 0.5, the DM is said to exhibit ambiguity-aversion, whereas he exhibits an ambiguity-seeking behaviour if α is inferior to 0.5. In cases of pure risk, the worst-case and best-case probability distributions coincide, and this intuitive definition matches the traditional definition of an expected cost.

To present more rigorously this criterion, let's assume $\mathcal{S}$ is the set of future states of the world. Ambiguity is represented by a set $\mathcal{C}$ of possible probability distributions over these states of the world. Let's consider $\mathbf{A}$, a set of (finitely-valued) mappings from $\mathcal{S}$ to $\mathbb{R}$, that we will refer to as *acts*, or *projects*. In other words, one can think of this representation as the assessment of the welfare gains or losses brought about by a project in each state of the world. An α-maxmin social-planner is then characterized by a utility function u and a coefficient of *ambiguity aversion* α, such that her preferences among acts can be represented by an index I, defined by equation 1:

$$\forall \mathbf{a} \in \mathbf{A},\ I(\mathbf{a}) = \alpha \min_{P \in \mathcal{C}} \int_S u(\mathbf{a})\, dP + (1-\alpha) \max_{P \in \mathcal{C}} \int_S u(\mathbf{a})\, dP. \tag{1}$$

This index is said to represent the preferences of the social planner if she prefers project $\mathbf{a}$ to project $\mathbf{b}$ if and only if $I(\mathbf{a})$ is superior to $I(\mathbf{b})$. This representation thus suggests that preferences among different projects, or choices among competing policies, can be determined by calculating the index I associated with each alternative, and then by choosing the project that has the largest index.

The normative rationale for the use of this decision criterion is nested in its axiomatic foundation. In other words, Ghirardato et al. (2004) showed that the α-maxmin criterion is equivalent to a set of six logical axioms. The strength of such a foundation is that a decision-maker that would want to abide by the set of axioms should feel compelled to using the equivalent decision criterion. In the following paragraphs, we are going to illustrate the implications of some of these axioms in terms of energy projects. We present these axioms in a more thorough way in the appendices.

The first two axioms state that one's preferences ought to be non-trivial, complete and transitive. This means that there must exist two projects among which the public evaluator is not indifferent; that any two projects must be either ranked or indifferent according to the evaluator's preferences; and that if project A is preferred to project B, and B is preferred to project C, then A ought to be preferred to C. The third axiom requires preferences to be monotonic. In other words, if project A scores better than project B in every future state of the world then A ought to be preferred to B. These are very usual axioms, which seem warranted in most decision problems. One critique that may be made to the monotonicity rule is that it may produce some unfair outcomes, according to how the score of a project in a particular state of the world is computed. In this paper, we assume that conditionally on the realization of a state of the world, the welfare losses associated with the analysed projects can be measured in monetary terms, and that this measure of welfare loss is somehow fair.

The last three axioms may seem less intuitive. Most importantly, preferences are required to be "certainty independent" (or C-independent). This axiom states that unrisky projects cannot hedge risky ones. In other words, it allows for the combination of several risky actions together in order to hedge the various risks entailed by each action taken separately, in order to reduce the overall level of risk entailed by the joint actions. The last two axioms are less appealing and essentially technical, their description is performed in the appendices.

2.2.2 Discussion of alternative approaches

Other decision rules exist in the same literature of decision-making under ambiguity. Namely, Gilboa and Schmeidler (1989) proposed the maxmin criterion, which coincides with the criterion we chose when $\alpha = 1$; and Klibanoff et al. (2005) proposed a *smooth* model of decision-making under ambiguity, which assigns a weight to every prior probability distribution in $\mathcal{C}$, instead of assigning positive mass to only the best and worst priors. We chose to use the α-maxmin criterion for the following reasons. First, it allows for various levels of attitudes towards ambiguity, while the maxmin criterion is based on a built-in ambiguity-aversion. Second, to compute the expected cost of some project, one only needs to specify the best case and worst case probability distributions that may apply to this particular project. This makes the criterion more tractable, and less informationally demanding than the smooth model of Klibanoff et al. (2005)[7].

[7] The smooth model of decision-making under ambiguity is informationally demanding as it requires the evaluator to come up with some *second-order* probability, which assigns each prior probability distribution in $\mathcal{C}$ with a weight. This "prior over the priors" can be interpreted as the probability of a prior to be the "true" prior associated with the project.

Two other branches of the theoretical economics literature also deal with social decisions made under risk, and ought to be mentioned here. The first one focuses on social choices made under risk: e.g. it tries to aggregate individual preferences into social choice rules that allow a public evaluator to determine the most desirable projects or policies. The main drawback of this literature is that it assumes that the probability distribution associated with the social risks is known to the decision maker. The latest contributions to this literature and its main references can be found in Fleurbaey (2010). The second branch concerns the aggregation of conflicting opinions of experts, which can be seen as some form of ambiguity. It was developed for example by Gajdos et al. (2008) (among others), whose results were applied by Crès et al. (2011) to the climate change example. As a matter of fact, the prescriptions found in these articles are similar to ours, as they suggest to decision maker willing to aggregate conflicting opinions to perform some weighted sum of the different opinions.

Finally, it can be interesting to notice that the method we propose is similar to the work of Henry and Henry (2002) who tried to formalize the precautionary principle. They studied the case of a social planner trying to maximize the utility of a representative agent when the different social alternatives available are subject to *scientific ambiguity* , which they define as the existence of several sources of scientific knowledge leading to different probability distributions over some states of the world. For instance, in our decision problem, Bayesian statistics and probabilistic safety analyses do not lead to the same probability of nuclear accidents. In addition, they show that under ambiguity, non-precautionary decision-making, e.g. decision-making based on Savage (1954) subjective-expected utility preferences among unambiguous acts, is suboptimal. As a consequence, they advocate the use of some decision criteria that exhibit ambiguity-aversion[8] and allow to consider ambiguous acts, as they embody rather well the guidelines of the precautionary principle, and can lead to increased levels of welfare. Our paper follows their guidelines as long as the coefficient α representing ambiguity attitude exhibits ambiguity-aversion. In this case, the existence of multiple probabilities associated with some events is accounted for by an increased level of pessimism.

3 An application to nuclear power accidents

We now propose to apply the α-maxmin decision criterion to the calculation of the expected cost of nuclear accidents. Following the definition of the criterion proposed in section 2, we will consider a function $\mathbf{a}$ which assigns to any state of the world $s \in \mathcal{S}$ the monetary valuation of the nuclear accident associated with this state, and calculate the value of the α-maxmin index $I(\mathbf{a})$ associated with this

[8] In their case, they use two decision criteria: the maxmin expected utility criterion of Gilboa and Schmeidler (1989), and the Choquet Expected Utility (CEU) criterion of (Schmeidler, 1989).

function. The first step of this process will consist in defining $\mathbf{a}$ as a function that assigns 0 to most states of the world and some monetary damage to the states that do represent a nuclear accident. We will then elicit the probability measures $P \in \mathcal{C}$ associated with different nuclear events, in a best-case and worst-case scenarios. We will finally present the results of the numerical application, perform some sensitivity tests, and relax some important hypotheses.

3.1 Nuclear accidents, damage and probabilities

3.1.1 Types of accidents

In 1990, the IAEA developed the International Nuclear and Radiological Event Scale (INES), which consists in a numerical rating of nuclear accidents on seven levels. Nuclear events are assessed on the basis of their outcomes, which are evaluated with respect to three criteria: people and the environment, radiological barriers and control, and defence-in-depth. The difference between two steps of INES is approximately a tenfold increase in the consequences of the event. Events ranked from level 1 to 3 are called nuclear incidents. The associated releases of radiological materials to the environment are necessarily below the acceptable thresholds, but on-site consequences can be serious. Events that are ranked above level 4 are called accidents. These levels entail severe on-site damage and radioactive releases in the environment that are at least as high as prescribed limits. Table 1 lists the numbers of observed accidents per INES level. It is taken from Cochran and McKinzie (2011).

Table 1: Observations of nuclear accidents since 1990 sorted by INES level

INES level	4	5	6	7
Number of observation	13	5	1	2

Following the choices made by IRSN (2013); Rabl and Rabl (2013) and D'Haeseleer (2013), the different nuclear accidents considered in this application are taken from a classification adopted by the NRC in 2009. We distinguish two categories of events. The first category will be called core-damage accidents (CDA). It is characterized by the heat-up of the reactor core up to the anticipation of advanced oxidation and severe fuel damage, involving a share of the core that would result in public health effects if released. Some of these core-damage accidents might develop further and become the second kind of accidents, called large-release accidents (LRA), which will embody large nuclear accidents as they are usually perceived. These accidents involve some unmitigated releases of airborne fission products to the environment (NRC, 2009). Historically, we can say that CDA refers to accidents similar to the one which occurred in Three-Mile Island, while LRA describes accidents similar to the ones which occurred in Fukushima-Daiichi or Chernobyl. Figure 2 summarizes the structure of these events, and their associated damage: D_{CDA} and D_{LRA}.

Figure 2: A simplified representation of nuclear accidents

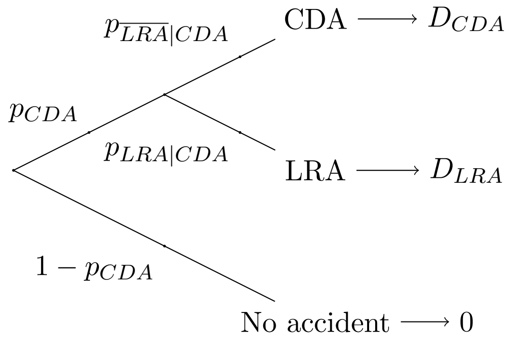

3.1.2 Nuclear damage

As we assumed in section 2 that the valuation of nuclear damage was out of the scope of this study, we use the existing literature to provide an estimate of D_{CDA} and D_{LRA}. Two categories of damage can be identified: those that are explicitly paid for by either the operator of the power station or by the Government once the liability threshold of the operator is exceeded; and those that are not explicitly paid for, such as the value associated with the lost use of environmental amenities. We believe that the scope of damage assessment that ought to be included in the decision process strongly depends on the nature of the decision-maker. For instance, while a government interested in setting its energy policy will have to consider the risks of large macroeconomic shocks caused by nuclear accidents, it might not be descriptively accurate to assume that a firm or a private investor would do so. In the following, we thus exclude the macroeconomic shocks caused by nuclear accidents and relax the assumption at the end of the section.

The estimation of the cost of a large-release accident was recently provided by France's nuclear safety technical body (IRSN, 2013). Their estimation for an accident taking place in France is €$_{2014}$ 449 billion. The figure provided by this study includes several components which are described further in the appendices. Two of these components, the *image cost*and *fleet cost*, are very specific to the French case. The fleet cost represents the expected regulatory upgrades required in the aftermath of a nuclear disaster. It is thus particularly high in France, because of its 58 operating reactors. The image cost represents the impact of a macroeconomic shock induced by a nuclear accident. For now, we subtract these costs from the estimation of total damage. Hence, our conservative estimation for the nuclear damage is €$_{2014}$ 180 billion for an LRA. Our conservative assumption for the cost of a core damage accident is the estimation by Sovacool (2008) of the cost of the Three-Mile Island accident, e.g. €$_{2014}$ 2.6 billion. This estimation includes the cost of the lost reactor, and tries to capture the consequences of the local panic that followed the TMI accident. Table 2 summarizes these damage estimates[9]. It is important to point out that we made a very strong assumption: damage valuations are assumed to be prone to no uncertainties; and ambiguity only affects the probabilities over events.

Table 2: The damage caused by nuclear accidents, in billion euros

Accident type	Damage estimate
Core Damage	2.6
Large Release	180

3.1.3 Probabilities of accidents

Given the framework described in section 2, we need to elicit the two prior probability distributions $P_{worst-case}$ and $P_{best-case}$ that yield the lowest and the highest values of $\int_S \mathbf{a}\, dP$, where $\mathbf{a}$ is a function that assigns a monetary value to each state of the world associated with a particular nuclear accident. The next paragraphs thus intend to construct the probabilities associated with the event-tree presented on figure 2 on page 13. These priors are summarized in table 3.

Probabilistic safety assessments are used to build the "best-case" prior. Indeed, PSAs have been used to derive total probabilities of accidents, even though they were not specifically designed for it. As they assume perfect compliance with safety standards and cannot foresee every chain of events that could lead to a nuclear accident, the figures obtained by PSAs are underestimation of the probabilities. They are thus a good candidate for the "best-case" prior.

As we restrict our study to the case of new builds, which have new designs (such as AREVA's EPR or Westinghouse's AP-1000 reactors), past observations of accidents may seem irrelevant for the characterization of their safety. On the other hand, we have seen that perceptions of probabilities were prone to an upward bias, and thus are a good candidate for the "worst-case" prior. Yet, there are no estimations of perceived probabilities of nuclear accidents. Therefore, we take the observed frequency of accidents in existing nuclear reactors as a proxy for public opinion. The rationale for this hypothesis is the post-Fukushima reactions of the media, which identified the observed frequency of accident with their future probability. As this assumption implies that new reactors and new technological designs are subject to some probability of accident calculated on the basis of the operation of older reactors, the prior based on this hypothesis will be a good candidate for the "worst-case" prior.

In the following, the probability p_{CDA} of a CDA will refer to the probability of occurrence of CDAs, while the probability of a LRA, p_{LRA}, will refer to the product of the probability of occurrence of a

[9] In previous version of this paper, we tried several pairs of damage estimations based on the recent studies dedicate to the assessments of nuclear damage, which are reviewed in Appendix 2.

CDA by the probability that a CDA evolves into a LRA, conditionally on the occurrence of a CDA. With respect to the notations introduced on figure 2 on page 13, we have $p_{LRA} = p_{CDA} \times p_{LRA|CDA}$.

The first prior is derived from AREVA's PSAs concerning its EPR project at Hinkley Point, we thus consider that the probability of a core damage accident is 10^{-6} per *reactor.year* and that the probability of a LRA is 10^{-7} per *reactor.year*. The second prior is more pessimistic. We consider a CDA probability of 1.10^{-3} per *reactor.year*, and a LRA probability of 1.10^{-4} per *reactor.year*. The latter figures were proposed by Rabl and Rabl (2013), and are the highest probabilities of nuclear accidents found in the literature review carried out by D'Haeseleer (2013)[10].

Table 3: The priors, expressed per *reactor.years*

Accident type	Best case prior	Worst case prior
CDA	10^{-6}	10^{-3}
LRA	10^{-7}	10^{-4}

A noticeable feature of table 3 is that the best-case and worst-case priors exhibit the same conditional probability of evolution of a CDA into a LRA. More precisely, in both cases, out of ten core damage accidents, only one further evolves into a large release accident. This assumption is supported by the existing literature, reviewed by D'Haeseleer (2013). Moreover, according to Cochran and McKinzie (2011), 18 nuclear events have been reported on the INES level 4-5, and 2 events were reported on level 7, which is roughly consistent with our hypothesis[11].

3.2 The expected cost of nuclear accidents

3.2.1 Parameters of the framework

Ambiguity attitude

The attitude of the decision-maker towards ambiguity is represented in our framework by the parameter α, which can take any values between 0 and 1. With $\alpha = 0$, we assume that the decision-maker behaves according to the *most optimistic* prior. She bases her decisions on the best-case scenario. Conversely, with $\alpha = 1$, she makes her decisions with respect to the *most pessimistic* prior or according to the worst-case scenario. When $\alpha = 0.5$, The decision-maker gives the same weights to both

[10]We are intentionally using the words "optimistic" and "pessimistic" rather than "conservative". Indeed, both the optimistic and pessimistic priors are conservative. "Optimistic" and "pessimistic" refer to the source of knowledge from which the priors are built. "Conservative" refers to the degree of precaution taken in the choice of the numbers in the priors. As an example, the optimistic prior is optimistic because it is based on AREVA's PSA study for the EPR, a source of information likely to be biased downward. It is also conservative in the sense that we chose to use the target probability of the PSA in our prior (e.g. 1.10^{-7}) rather than AREVA's estimation of the probability of a large release accident, which is 2.10^{-8} (see table 1 on page 6).

[11]The only level 6 event reported in table 1 did not occur in a commercial power station.

scenarios. In the following application, we will present the expected cost of nuclear accidents as a function of α.

Load factor and nominal power

To compute the expected cost of nuclear accidents in euros per megawatt.hour (€/MWh), we need to know the electrical output of the power plant subject to the risk of accident. As most countries use nuclear power as a base-load technology, we hypothesize that the mean yearly load factor of the studied reactor will be 85%. This assumption is rather conservative, since the average load factors of nuclear power stations in the US was 91.8% in 2014, according to the NEI[12]. Load factors will be referred to by the letter ρ. The nominal power of the reactor considered is that of the European Pressurized Reactor, e.g. 1650 MWe. This nominal power is referred to in the forthcoming equations by the symbol P_{EPR}. The total production of the reactor over a year is calculated by multiplying the nominal power by load factor and by the number of hours in a year.

3.2.2 Numerical application

Given the priors expressed in table 3, and the two estimations of total damage expressed in table 2, we can calculate the α-maxmin expected cost of nuclear accidents. The expected cost relative to either of our prior probabilities can be calculated in the following way:

$$\int_{\mathcal{S}} \mathbf{a} dP = p_{LRA} \times D_{LRA} + (p_{CDA} - p_{LRA}) \times D_{CDA}$$

The α-maxmin expected cost $EC(\alpha)$, normalized by the electrical output of a reactor, thus boils down to:

$$EC(\alpha) = \frac{\alpha \int_{\mathcal{S}} \mathbf{a} dP_{worstcase} + (1-\alpha) \int_{\mathcal{S}} \mathbf{a} dP_{bestcase}}{\rho * 365,25 * 24 * P_{EPR}}$$

The results of this calculation are plotted on figure 3. The most pessimistic result, e.g. $EC(\alpha = 1)$ remains inferior to 1.7€$/MWh$.

3.2.3 Sensitivity analysis and image costs

We performed a simple sensitivity analysis on the value of the expected cost of nuclear accident with $\alpha = 1$, e.g. the upper right end of figure 3. The sensitivity analysis was performed with respect to the values chosen for the probabilities and damage of CDAs and LRAs. Results are provided on figure 4, which presents the variations of $EC(\alpha = 1)$ with respect to variations of D_{CDA}, D_{LRA}, P_{CDA}

[12]Nuclear Energy Institute

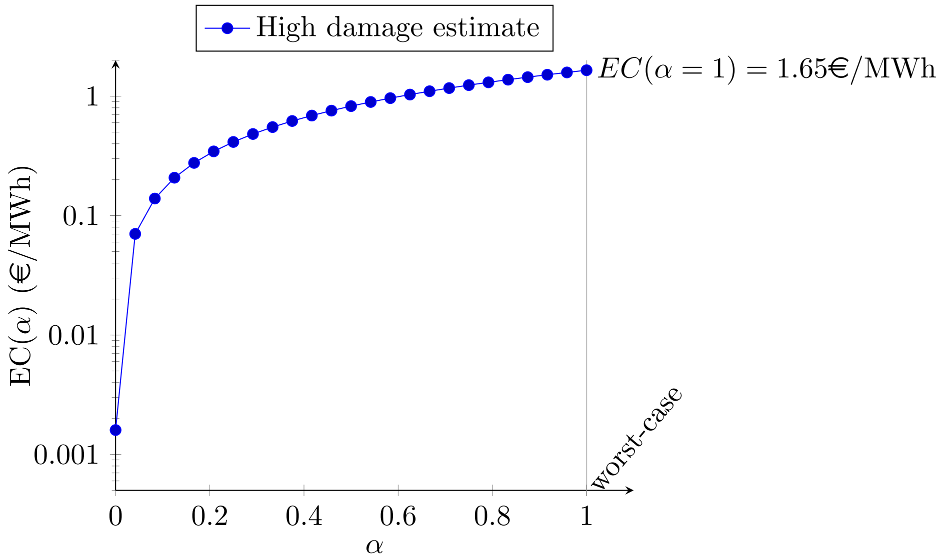

Figure 3: expected cost in €/MWh as a function of α

and P_{LRA} (variations are indicated in percentage of the values chosen in the application). Given the linearity in probabilities and damage, the sensitivity is linear. Yet, this sensitivity test confirms that, with the priors chosen and the hypotheses made on nuclear damage, the expected cost is more sensitive to the damage and probabilities of LRAs than to those of CDAs. This result is comforting, as it means that decisions regarding energy alternatives are more impacted by the possibility of witnessing rare but dreadful nuclear accidents than by the possibility of destroying reactors that do not release radioactive materials into the environment. It also appears that the result is more sensitive to the damage than to the probabilities. Therefore, we will discuss our choices of damage estimates in the next section.

The next sensitivity test consists in releasing the assumption according to which the image costs of nuclear accidents are null. We thus increase the damage estimates listed in table 2 on page 14 to account for the impact of a macroeconomic shock induced by a nuclear accident. The evaluation of the cost of such a macroeconomic shock is taken from the IRSN (2013) study, and the new damage estimates are gathered in table 4, where Δ represents the ratio between the new and the former estimates. It appears that image costs account for half of the damage caused by LRAs, and for a very large share of the damage caused by CDAs. Figure 5 presents the new results: under the most stringent hypotheses, the expected cost is approximately 7 €/MWh. We would like to stress that the accuracy of this numerical result is certainly not general, as the consequences of a macroeconomic shocks do not satisfy the hypothesis of no uncertainty on damage. Moreover, macroeconomic shocks are necessarily dependant on the location of the nuclear accident, and the very high value taken by the image damage in the IRSN study are justified by the strong reliance of the French economy on its

Figure 4: Sensitivity study: expected cost versus damage and probabilities

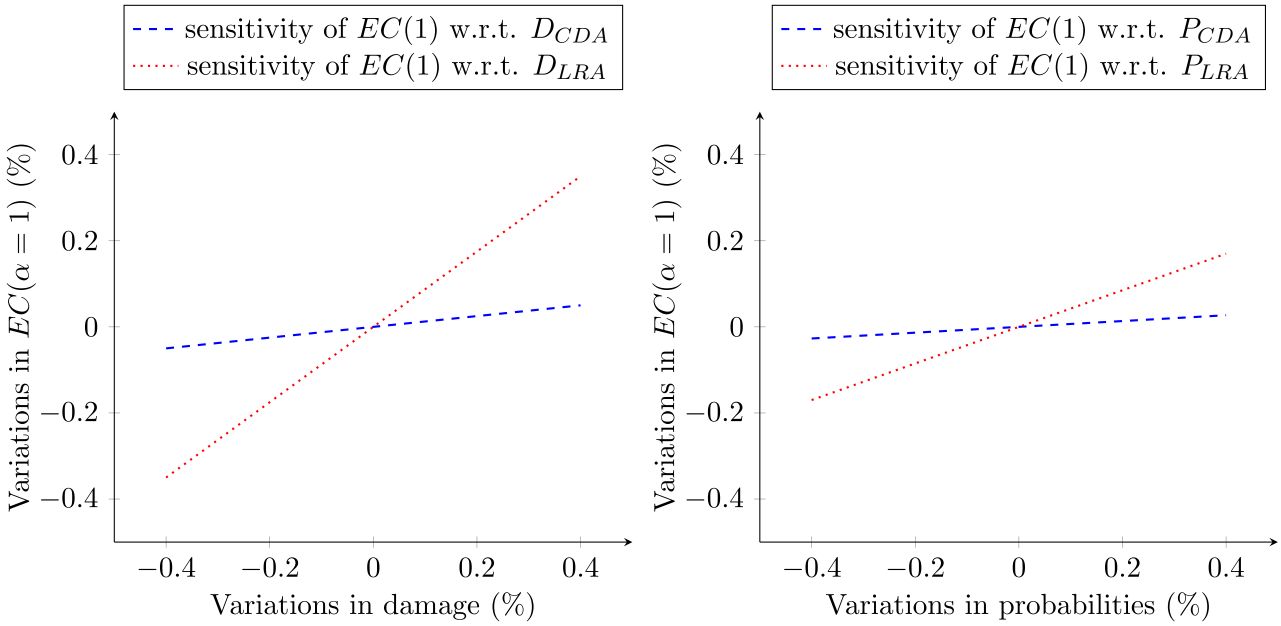

tourism and agriculture.

Table 4: Damage estimation including image costs, and variations with respect to the previous hypotheses

Accident type	New damage estimate (in billion €)	Δ
CDA	52	x20
LRA	361	x2

Finally, we can also relax the “no fleet cost” hypothesis, by expressing the variation of the α-maxmin expected cost when other reactors have to be upgraded or shut down as a consequence of the accident. To do so, we consider the estimation of fleet costs for France made by the IRSN (2013). They assume that this fleet cost is of approximately $€_{2014}$ 90 billion. France counts 58 operating nuclear reactors. If we assume that an accident destroys one of these reactors, and that upgrade costs are equally split among remaining reactors, it means that each reactor incurs a cost of $€_{2014}$ 1.6 billion in case a LRA occurs. We can thus compute the expected additional cost incurred per additional reactor by applying the α-maxmin valuation rule to a function that associates $€_{2014}$ 1.6 billion to states of the world that correspond to large-release accidents, and 0 to other states of the world. For instance, with $\alpha = 1$, we can say that the expected cost of nuclear accidents is 1.7 €/MWh plus 0.013 €/MWh per reactor in the fleet.[13] We thus see that for most countries which do not rely extensively on nuclear

[13] 7 €/MWh plus 0.013 €/MWh per reactor in the fleet if macroeconomic shocks are included.

Figure 5: expected costs with image costs

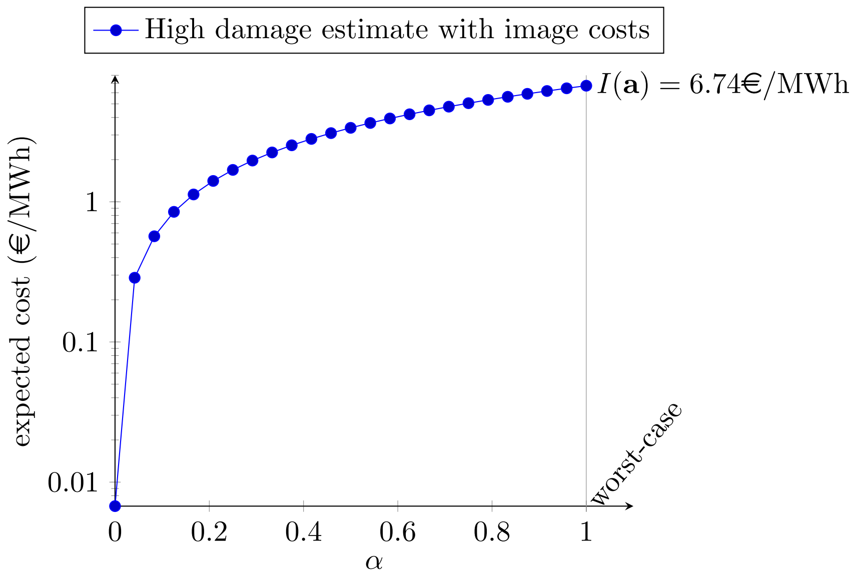

power, the fleet effect should not be an important driver of decisions related to the use of nuclear power. In countries in which a large fleet exists, such as the United-States, France, China, Japan or Korea, this effect can account for a significant part of the expected cost of future nuclear accidents in new builds.

4 Discussion and policy implications

This paper develops a method for the calculation of the expected costs of rare and catastrophic energy-related accidents, that takes into account the ambiguity that characterizes the probabilities of these events. Our method adapts a theoretical decision rule taken from the literature dedicated to decision-making under uncertainty. The expected cost provided by this method is no longer the result of the aggregation of the various types of damage incurred by society in the aftermath of an accident, but an index associated with any decision that may bring about a nuclear accident. The relevance of this index lies in its axiomatic foundation: a social-planner who would want her choices among energy projects or policies to be consistent with the axioms presented by Ghirardato et al. (2004), should evaluate rare disasters according to the rule we derived. We also apply this method in order to derive the expected cost of nuclear accidents in new builds. This expected cost is evaluated at 1.7 €/MWh, or 7 €/MWh when accounting for macroeconomic consequences.

We can now focus on the policy implications of these figures. First of all, the knowledge of the costs of nuclear power accidents is important for the four following general policy issues: the insurance and compensation of victims of nuclear accidents, the phase-out of ageing nuclear power plants, the determination of the optimal electrical mix of production technologies, and the choice of the appropriate level of safety investments[14]. Our method does not seem to shed light on the two former policy issues, as we haven't been dealing with individual damage undergone during a nuclear accident, and as our method does not concern currently operated power plants. Yet, the expected costs calculated in this paper bring some new light on the two latter issues.

Indeed, as our estimation is derived from a general decision criterion, the number it yields is well suited to be used as a comparative index between different alternatives that may or may not involve nuclear power reactors. In other words, when choosing among different energy scenarios, the risks associated with different technologies could be compared by referring to our figures, or by using this general method. Likewise, when setting safety standards for firms which are subject to ambiguous risks, our method provides a tractable way to measure the marginal expected costs and benefits of modifying existing safety standards, as long as these modifications can be associated with changes in the multiple probabilistic priors or with the damage that would be caused by an accident.

Quantitatively, the numbers presented in this study do not call for an immediate reconsideration of the use of nuclear power. Indeed, the levelized cost of energy generation (LCOE) provided by the AIE (2015) suggests that the LCOE of new nuclear power is approximately 90 €/MWh, which is more than an order of magnitude larger than our worse-case estimation. Besides, the other LCOE estimates provided by the AIE for diverse other technologies are such that an increase of 7€/MWh of the LCOE of new nuclear power would not change radically its comparative competitiveness when compared to either gas or coal plants, or with modern renewable sources of energy. The main reason for this policy implication is that the order of magnitude of the LCOEs of most modern technologies of production of energy varies between 50 €/MWh and 150 €/MWh, and that very few technologies are in such close competition that 7 €/MWh would change the balance between one or the other. Besides, it could be interesting to use the method developped in this paper to compare the expected costs of nuclear accidents with the expected costs associated with other rare energy disasters such as oil spills or dam failures.

Beyond its policy implications, it is possible to point out several shortcomings of our assessment. First, in the static framework that was presented in section 2, decisions are optimal *ex ante*, and with respect to the available information. Hence, the methodology proposed here is subject to unknown unknowns, which can be seen as events whose probabilities are null for all probability distributions within the set of possible priors. Besides, with a dynamic setting in which information could evolve

[14]One could add to these issues the choice of the localization of new power stations. We consider this question to be included in the choice of the appropriate level of safety investments.

over time, one could argue that it might be better *ex ante* to choose energy alternatives that allow one to adapt his choices to the acquisition of new information in the future. Such preferences have been studied by Kreps (1979), who modelled the behaviour of individuals characterized by preferences for flexibility: being uncertain about their future preferences, they would rather keep a combination of options than choose among them right away. This refinement is left for future research.

Finally, the hypothesis that ambiguity only characterizes the probabilities of occurrence of rare and dreadful events is another shortcoming of this method. Assessing the damage of such events is also a source of uncertainty, even though we believe that the figures we used as damage estimates in this study were derived in a tractable way and are conservative with respect to the existing literature. Nonetheless, the existence of multiple techniques for the valuation of non-monetary losses, and the large variety of damage caused by events such as nuclear accidents tend to suggest that an interesting path for future research would be to take into account this additional source of ambiguity in the assessment of the risks of energy disasters. Choices among acts characterized by ambiguous outcomes rather than by ambiguous probabilities have been discussed by Jaffrey and Jeleva (2009). To the best of our knowledge, there is no literature that addresses issues in which ambiguity characterizes both outcomes and probabilities.

Appendix 1: The logical foundation of our decision criterion

Let n be the number of available technologies dedicated to the production of energy. These technologies will be noted T_i, $i \in [\![1;n]\!]$. Assume that a decision-maker is interested in the determination of the optimal portfolio of these n technologies. A combination of the n technologies will be called an *energy alternative*. The set of all possible alternatives is $\mathcal{A}$. Elements of $\mathcal{A}$ are defined by: $\forall a \in \mathcal{A}, a = (a_i)_{i\in[\![1;n]\!]}$ such that $\sum_i a_i = 1$, in which a_i is to be interpreted as the share of technology i in the total production of energy desired by our decision-maker.

Let $\mathcal{S}$ be the set of future states of the world, and Σ a σ-algebra of subsets of $\mathcal{S}$, called events. Without loss of generality, we will consider potential future states of the world as combinations of consequences brought about by each available energy technology. We further assume that, conditionally on the realization of one particular state of the world, the welfare losses induced by the consequences of the use of the technologies included in the chosen energy alternative can be measured in monetary terms. This assumption means that the cost of any outcome of the use of a technology can be assessed *ex post*. The rationale for this assumption is the following: even though we acknowledge that the existing valuation methodologies dedicated to environmental or non-monetary welfare losses are imperfect, they nonetheless exist and allow one to provide assessments of the costs of past accidents. As we wish to tackle the decision problems raised by events which fail to representable by a classical probabilistic lottery, this hypothesis allows us to focus on the probabilistic side of the problem.

We thus suggest that any energy alternative can be described by an *ex-post* valuation function on $\mathcal{S}$ with values in $\mathbb{R}$, that associates each state of the world s with the sum of the monetary valuations of the consequences of the use of each energy production technology included in the given alternative. For instance, an alternative that would only involve T_1 would associate each state of the world s with the monetary value of the consequences brought about by the use of T_1 in state s. Likewise, an alternative that involves several technologies associates each state of the world with the weighted sum of the valuations of the consequences of each technology. This implies that each energy alternative a can be associated with a function $\mathbf{a} : \mathcal{S} \to \mathbb{R}$. Let $\mathbf{A}$ be the set of functions that represent some a in $\mathcal{A}$. By construction, an element of $\mathbf{A}$ corresponds to one and only one element of $\mathcal{A}$. Therefore, elements of $\mathbf{A}$ and $\mathcal{A}$ will be referred to as energy alternatives in an indifferent way in the remainder of this paper.

Mimicking the classical framework of decision theory, an alternative $\mathbf{a}$ that assigns the same element of $\mathbb{R}$ to each state of the world will be called a constant alternative. The set of constant alternatives will be noted $\mathbf{A}^c$. Moreover, for any $\lambda \in [0;1]$, the alternative $\lambda\mathbf{a} + (1-\lambda)\mathbf{b}$, is called a mixture, and designates the alternative whose i-th component is $\lambda a_i + (1-\lambda)b_i$. We also define *preferences* over $\mathbf{A}$ as a binary relation over elements of $\mathbf{A}$. These preferences characterize our decision-maker, and will be noted $\preceq$. Hence, $\mathbf{a} \preceq \mathbf{b}$ will be read: $\mathbf{b}$ is preferred to $\mathbf{a}$. Following the work of Ghirardato et al. (2004), we also define *unambiguous preferences*, $\preceq^*$, as follows:

$$\forall \mathbf{a}, \mathbf{b} \in \mathbf{A},\ \mathbf{a} \preceq^* \mathbf{b} \Leftrightarrow \forall \mathbf{c} \in \mathbf{A}, \forall \lambda \in [0,1], \lambda\mathbf{a} + (1-\lambda)\mathbf{c} \preceq \lambda\mathbf{b} + (1-\lambda)\mathbf{c}.$$

In other words, the preference among two alternatives is said to be *unambiguous* when it cannot be reverted by mixing the two alternatives with a third alternative. Given any $\mathbf{a}$ in $\mathbf{A}$, the set of unambiguous certainty equivalents of $\mathbf{a}$, $C^*(\mathbf{a})$, is defined as follows: $C^*(\mathbf{a}) = \{\mathbf{b} \in \mathbf{A}^c, \mathbf{a} \preceq^* \mathbf{b}$ and $\mathbf{b} \preceq^* \mathbf{a}\}$. This framework is almost identical to that of Ghirardato et al. (2004). The main difference is that Ghirardato's "space of *outcomes*", e.g. the image of $\mathcal{S}$ by functions in $\mathbf{A}$, is here assumed to be an interval of $\mathbb{R}$. This modification is a particular case of the GMM framework, which only requires outcomes to be a convex subset of a vector space. A similar simplification can be found in Chateauneuf et al. (1996).[15].

[15]Another difference is that Ghirardato's *acts* are here called *alternatives*

The decision problem is to identify within $\mathbf{A}$ the most desirable energy alternatives. This decision problem relates to the identification of the preferences of our decision-maker with respect to the energy alternatives in $\mathbf{A}$. Let's assume that our decision-maker wishes her decisions in terms of energy choices to be somehow rational. For instance, a usual rational postulate, or axiom, is transitivity: if alternative $\mathbf{a}$ is preferred to alternative $\mathbf{b}$ and $\mathbf{b}$ is preferred to a third alternative $\mathbf{c}$, then our decision-maker may want $\mathbf{a}$ to be preferred to $\mathbf{c}$. To establish our decision criterion, we assume that our decision-maker's preferences towards energy alternatives respect the following axioms, proposed by Ghirardato et al. (2004).

Axiom 1: Weak order. This axiom requires the preference relation between alternatives to be complete and transitive. Transitivity has been explained already. Completeness requires that given any two alternatives $\mathbf{a}, \mathbf{b}$ in $\mathbf{A}$, $\mathbf{a} \preceq \mathbf{b}$, or $\mathbf{b} \preceq \mathbf{a}$. Both statements can be true, in which case we write $\mathbf{a} \sim \mathbf{b}$, and say that alternatives $\mathbf{a}$ and $\mathbf{b}$ are equivalent.

Axiom 2: Certainty independence. This axiom states that preferences among a pair of alternatives cannot be hedged by a mixture with a constant alternative. More precisely, given any two alternatives $\mathbf{a}, \mathbf{b}$ in $\mathbf{A}$, such that $\mathbf{a} \preceq \mathbf{b}$, C-independence requires:

$$\forall \mathbf{c} \in \mathbf{A}^c, \forall \lambda \in [0, 1], \lambda \mathbf{a} + (1 - \lambda)\mathbf{c} \preceq \lambda \mathbf{b} + (1 - \lambda)\mathbf{c}.$$

The certainty independence axiom translates the idea of risk hedging. Indeed, it only requires independence over constant alternatives (alternatives that yield the same outcome whatever the future state of the world). Consequently, for mixtures of non-constant alternatives, uncertainties may add-up or hedge each other so as to change the preferences of the decision-maker. This is a weak form of the independence axiom proposed by Von Neumann and Morgenstern (1947), which requires that the preference among two alternatives remains unaltered by any mixture of these alternatives with any other alternative.

Axiom 3: Continuity. Let $\mathbf{a}$, $\mathbf{b}$ and $\mathbf{c}$ be three energy alternatives, such that $\mathbf{a} \prec \mathbf{b} \prec \mathbf{c}$. Then there exist λ_1, $\lambda_2 \in (0; 1)$ such that $\lambda_1 \mathbf{a} + (1 - \lambda_1)\mathbf{c} \prec \mathbf{b} \prec \lambda_2 \mathbf{a} + (1 - \lambda_2)\mathbf{c}$. Continuity is a technical axiom that enables the mathematical proof of the following representation proposition. It can nevertheless be interpreted as the assumption that no alternative is so bad (respectively so good), that no matter what mixture is made with it, this mixture is still not preferred (respectively preferred) to other alternatives.

Axiom 4: Monotonicity. If the outcome of an alternative $\mathbf{a}$ is greater than the outcome of another alternative $\mathbf{b}$ in all future states of the world, then $\mathbf{b} \preceq \mathbf{a}$.

Axiom 5: Non-degeneracy. There exist two alternatives such that one is strictly preferred to the other. This a technical axiom that ensures that the problem is not trivial, e.g. that not all alternatives are equivalent.

Axiom 6: Certainty equivalence. This axiom states that all the information required by a decision-maker to evaluate energy alternatives is contained within their sets of unambiguous certainty equivalents. More precisely, axiom 6 requires:

$$\forall \mathbf{a}, \mathbf{b} \in \mathbf{A}, C^*(\mathbf{a}) = C^*(\mathbf{b}) \Rightarrow \mathbf{a} \sim \mathbf{b}.$$

Hence, the result of Ghirardato et al. (2004) can be applied to the preference relation of our decision-maker. We use a weaker form of the theorem, that takes into account our modification of the structure of Ghirardato's outcome space X[16]. Therefore, $\preceq$ satisfies axioms 1 to 6 if and only if there exist a non-empty, weak* compact and convex set $\mathcal{C}$ of probabilities on Σ, and $\alpha \in [0,1]$ such that $\preceq$ is represented by the functional I, defined on $\mathbf{A}$ by:

$$\forall \mathbf{a} \in \mathbf{A},\ I(\mathbf{a}) = \alpha \min_{p \in \mathcal{C}} \int_S \mathbf{a}\, dP + (1-\alpha) \max_{p \in \mathcal{C}} \int_S \mathbf{a}\, dP.$$

Moreover, $\mathcal{C}$ is unique, and α is unique provided $\mathcal{C}$ is not a singleton. The statement "the functional I *represents* $\preceq$" means that $\forall \mathbf{a}, \mathbf{b} \in \mathbf{A}, \mathbf{a} \preceq \mathbf{b} \Leftrightarrow I(\mathbf{a}) \leq I(\mathbf{b})$.

Appendix 2: The assessment of nuclear power damage

The method proposed in this paper stands on the strong hypothesis that the valuation of damage caused by a rare disaster is always available and prone to no uncertainty. As this is obviously a very strong assumption, we present here the existing assessments of nuclear damage, in order to give further tractability to the numerical application of our method. The assessment of the damage associated with nuclear accidents in commercial power stations has been mostly motivated by previous catastrophes. It was first assessed in the United-States after the Three-Mile Island accident (Rasmussen, 1975; Nuclear Regulatory Commission, 1982). A series of studies were dedicated to the the consequences

[16] Provided the outcome space X is included in $\mathbb{R}$, and as Ghirardato's non-constant affine utility function u is obtained independently from $\mathcal{C}$, and is defined up to a positive, affine transformation, we can, without loss of generality, assume that u is the identity function.

of the Chernobyl accident (Ottinger et al., 1990; Hohmeyer, 1990; Ewers and Rennings, 1991, 1992; CEPN, 1995). Finally, some studies have started to assess the cost of the Fukushima-Daiichi accident in 2011 (Versicherungsforen, 2011; Rabl and Rabl, 2013). More recently, France's nuclear safety technical body, the IRSN[17], published an assessment of the cost of nuclear accidents (IRSN, 2013). Table 5 on page 25 presents these studies.

Table 5: Existing assessments of the damage caused by major nuclear accidents.

Sources	Year	Health damage	Agri-culture	Production losses	On-site costs	Image costs	Fleet costs	Total (G€$_{2014}$)
WASH 1400	1975	†	†	†	*	*	*	63
CRAC-2	1982	†	†	†	*	*	*	657
Hohmeyer	1988	1,370	*	*	*	*	*	2,174
Ottinger	1990	629	38	*	*	*	*	989
Ewers *et al*	1991	2,740	38	828	*	*	*	5,179
Ewers *et al*	1992	7,815.6	307.4	179.1	*	*	*	11,648
ExternE	1995	74.3	*	37.9	*	*	*	149
Eeckhoudt	2000	10.85	6.162	0.098	*	*	*	428
GREF	2011	†	†	†	*	*	*	6,097
Rabl-Low	2012	10	5	100	50	*	*	167
Rabl-Central	2012	18.8	7.5	250	78	*	*	359
Rabl-High	2012	50	50	1,000	290	*	*	1,409
IRSN-severe	2013	0	9	11	10	50	44	125
IRSN-major	2013	27	14	110	28	180	88	449

Damage categories are expressed in current euros. Total is expressed in constant euros.

Damage Categories: † : details not provided * : not accounted for

Table 5 shows high variations. How can one assess the cost of nuclear accidents at approximately €63 billion (Rasmussen, 1975), while others derive a cost of more than a trillion euros? (Ewers and Rennings, 1992) Likewise, even though it only assesses health, food and production costs, the Versicherungsforen (2011) study calculates a total cost ten times superior to the IRSN (2013) figure, which accounts for a larger panel of consequences. The following paragraphs intend to show that these variations arise from both the assessment of the physical consequences of the accidents and the methodologies dedicated to the calculation of the monetary equivalents of these physical consequences.

Assessing physical consequences

The consequences of a nuclear accident are numerous and intricate. An accident can have on-site consequences, such as casualties, highly-irradiated workers or material losses in adjacent reactors. It can also cause various off-site consequences, such as the release of radioactive materials in the atmosphere, the contamination of lands or cattle, or a reduced attractiveness for tourism. As different

[17]IRSN stands for Institut de Radioprotection et de Sûreté Nucléaire

as they may be, these numerous consequences have to be assessed in order to derive the cost of the accident.

The divergence in the assessment of the consequences is twofold. First, all studies do not assess the same range of consequences. Some studies argue that health effects dominate all other effects (Versicherungsforen, 2011; Hohmeyer, 1990; Ottinger et al., 1990). They thus focus on the collective absorbed dose and neglect other consequences. Other studies focus on a wider panel of effects, such as land exclusion, or country image effects (tourism, exportations of foodstuffs), and fleet effects (regulatory upgrades after the accident). The studies on which we based the calculations of section 2 belong to this category.

Second, studies also differ in their assessment strategies. Radiological consequences can be modelled by dedicated programs (MACCS, COSYMA... (Jones et al., 1996)) that rely on level-three probabilistic safety assessments. Consequences have also been assessed by adapting the figures derived from past catastrophes: most studies performed in the early nineties were based on Chernobyl's figures. They found particularly high values for the total cost of the accident (Hohmeyer, 1990; Ottinger et al., 1990; Ewers and Rennings, 1992). This raises an important question. Can we rely today on studies that have assessed future accidents on the basis of the consequences of past catastrophes? The answer is that we cannot. Relying on past figures fails to account for the learning from past consequences, the enhancement of safety standards, and the progress in available mitigation technologies.

The economic valuation of the consequences

Once the consequences of a nuclear accident have been assessed in physical terms (number of casualties, absorbed dose of radiations, area of polluted lands...), they have to be given a monetary value. The monetary equivalent of some consequences, such as material losses, will be easily derived. Unfortunately, limited knowledge impedes the proper estimation of other physical consequences, such as environmental or health damage.

To give a few examples of this limited knowledge, we do not know precisely the effect of the exposure to low doses of radiations on the probabilities of contracting cancers or hereditary diseases (Kathren, 1996). The consequences on food are also uncertain since the population can react to food-bans by boycotting healthy products. The harm caused by nuclear countermeasures, such as the psychological distress of people due to relocations, is also hard to assess.

Environmental losses can be assessed by the evaluation of individuals willingness-to-pay (WTP) to avoid the losses. Two families of methods have been developed to assess this WTP: the revealed-preference methods and the stated-preference methods. Revealed-preference methods can be used in order to evaluate environmental losses, but they rely on past behaviours and require data, they are thus hard to apply to nuclear accidents (Knetsch and Sinden, 1984; Hanemann, 1991; Adamowicz et al., 1994). Stated-preference methods are based on surveys that try to elicit the willingness-to-pay of people to restore the environment. Namely, the "contingent-valuation method" is often used to value the environmental consequences of rare disasters.

Regarding health costs, the economic value of fatalities or impairments is usually assessed with the human capital method. This method consists in an assessment of the number of lost years of production which are multiplied by the average yearly production of a human being. It usually yields values of lost lives of approximately two to six millions euros. Other methods have been developed and propose different ways of calculating those health costs (Viscusi, 1993; Kip Viscusi and Aldy, 2003; Fautrel et al., 2007; Van den Hout, 2010).

This variety of methodologies is responsible for some of the discrepancies observed in table 5. First, a given consequence can be assessed with different methods which may yield different monetary values. Second, even when different studies assess the same consequence with the same method, results can still be different. As an example, Hohmeyer (1990) and Ottinger et al. (1990) both assess the value of lost life with the human-capital method. Yet, they respectively use a $1 million value and a $4 million value.

To conclude this appendix, we believe that the figures we used as damage estimates in this study were derived in a tractable way and are conservative with respect to the existing literature. Nonetheless, the existence of multiple techniques for the valuation of non-monetary losses, and the large variety of damage caused by nuclear accidents tend to suggest that an interesting path for future research would be to take into account this additional source of ambiguity in the assessment of the risks of energy disasters. Choices among acts characterized by ambiguous outcomes rather than by ambiguous probabilities have been discussed by Jaffrey and Jeleva (2009). To the best of our knowledge, there is no theoretical literature that addresses issues where ambiguity characterizes both outcomes and probabilities.

References

Adamowicz, W., Louviere, J., and Williams, M. (1994). Combining revealed and stated preference methods for valuing environmental amenities. *Journal of Environmental Economics and Management*, pages 271–292.

AIE (2015). Levelized cost and levelized avoided cost of new generation resources in the annual energy outlook 2015. Technical report, U.S. Energy Information Administration.

Azar, C. and Lindgren, K. (2003). Catastrophic events and stochastic cost-benefit analysis of climate change. *Climatic Change*, 56(3):245–255.

Bewley, T. F. (2003). Knightian decision theory, part i. *Decisions in Economics and Finance*, 25:79–110.

Bureau of Reclamation (2015). Best practices in dam and levee safety risk analysis. Technical report, U.S. Department of Interior and U.S. Army Corps of Engineers. http://www.usbr.gov/ssle/damsafety/risk/methodology.html.

CEPN (1995). Externe : Externalities of energy - volume 5 : Nuclear. Technical report, Luxembourg.

Chateauneuf, A., Kast, R., and Lapied, A. (1996). Choquet pricing for financial markets with frictions. *Mathematical Finance*, 6(3):323–330.

Cochran, T. B. and McKinzie, M. G. (2011). World federation of scientists' international seminars on planetary emergencies. Erice. Ettore Majorana Centre.

Crès, H., Gilboa, I., and Vieille, N. (2011). Aggregation of multiple prior opinions. *Journal of Economic Theory*, (146):2563–2582.

D'Haeseleer, W. D. (2013). Synthesis on the economics of nuclear energy. Technical report, DG Energy.

Dow, J. and Werlang, S. R. d. C. (1992). Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica*, pages 197–204.

Eeckhoudt, L., Schieber, C., and Schneider, T. (2000). Risk aversion and the external cost of a nuclear accident. *Journal of Environmental Management*, pages 109–117.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, 75:643–669.

Epstein, L. G. and Schneider, M. (2008). Ambiguity, information quality and asset pricing. *The Journal of Finance*, 63(1):197–228.

Epstein, L. G. and Wang, T. (1994). Intertemporal asset pricing under knightian uncertainty. *Econometrica: Journal of the Econometric Society*, pages 283–322.

Escobar Rangel, L. and Lévêque, F. (2014). How fukushima dai-ichi core meltdown changed the probability of nuclear accidents ? *Safety Science*, 64:90–98.

Ewers, H.-J. and Rennings, K. (1991). Die volkswirtschatlichen kosten eine super-gau's in biblis. *Zeitschrift für Umweltpolitik und Umweltrecht*, pages 379–396.

Ewers, H.-J. and Rennings, K. (1992). *Abschätzung der Schäden durch einen sogennanten Super-GAU*. Prognos, Bâle.

Fautrel, B., Clarke, A. E., Guillemin, F., Adam, V., St-Pierre, Y., Panaritis, T., Fortin, P. R., Menard, H. A., Donaldson, C., and Penrod, J. R. (2007). Costs of rheumatoid arthritis: new estimates from the human-capital method and comparison to the willingness-to-pay method. *Medical Decision Making*, 27:138–150.

Fischhoff, B., Slovic, P., and Lichtenstein, S. (1978). How safe is safe enough? a psychometric study of attitudes towards technological risks and benefits. *Policy Sciences*, pages 127–152.

Fleurbaey, M. (2010). Assessing risky social situations. *Journal of Political Economy*, 118(4):649–680.

Gajdos, T., Tallon, J.-M., and Vergneaud, J.-C. (2008). Representation and aggregation of preferences under uncertainty. *Journal of Economic Theory*, (141):68–99.

Ghirardato, P., Maccheroni, F., and Marinacci, M. (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory*, 118:133–173.

Gilboa, I. (2004). *Uncertainty in Economic Theory - Essays in honor of David Schmeidler's 65th birthday*. Routledge - Taylor and Francis Group, London.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18:141–153.

Hanemann, W. M. (1991). Willingness to pay and willingness to accept: How much can they differ? *The American Economic Review*, 81:635–647.

Hanley, N. and Spash, C. L. (1993). *Cost Benefit Analysis and the Environment*. Edward Elgar Publishing Limited, Chelthenham.

Henry, C. and Henry, M. (2002). Formalization and applications of the precautionary principles. Columbia Discussion Papers Series.

Hofert, M. and Wüthrich, M. V. (2011). Statistical review of nuclear power accidents. *Asia-Pacific Journal of Risk and Insurance*, 7:1–13.

Hohmeyer, O. (1990). Latest results of the international discussion on the social costs of energy - how does wind compare today? Technical report, Madrid.

Hurwicz, L. (1951). A class of criteria for decision-making under ignorance. *Cowles Commission Paper*.

IAEA (2006). *Chernobyl's Legacy: Health Environmental and Socio-Economic Impacts and Recommendations to the Governments of Belarus, the Russian Federation and Ukraine*. The Chernobyl Forum: 2003-2005, Vienna, second revised edition edition.

IAEA (2010). Development and application of level 1 probabilistic safety assessment for nuclear power plants. Technical report.

IRSN (2013). Méthodologie appliquée par l'IRSN pour l'estimation des coûts d'accidents nucléaires en France. Technical report.

Jaffrey, J.-Y. and Jeleva, M. (2009). How to deal with partially analyzable acts? *Theory and Decision*, 71(1):129–149.

Jones, J., Mansfield, P. A., and Hasemann, I. (1996). Pc cosyma version 2: An accident consequence assessment package for use on a pc. Technical report.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Kathren, R. L. (1996). Pathway to a paradigm: The linear nonthreshold dose-response model in historical context. *Health Physics*.

Kip Viscusi, W. and Aldy, J. E. (2003). The value of a statistical life: A critical review of market estimates throughout the world. *The Journal of Risk and Uncertainty*, 27:5–76.

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

Knetsch, J. L. and Sinden, J. A. (1984). Willingness to pay and compensation demanded: Experimental evidence of an unexpected disparity in measures of value. *The Quarterly Journal of Economics*, pages 507–521.

Knight, F. H. (1921). *Risk, uncertainty and profit.* Hart, Schaffner and Marx, New York.

Kreps, D. M. (1979). Preference for flexibility. *Econometrica*, 47(3):565–577.

Lévêque, F. (2013). *Nucléaire On/Off - Analyse économique d'un pari.* Dunod, Paris.

NRC (2009). An approach for determining the technical adequacy of probabilistic risk assessment results for risk-informed activities. Technical report, Office of Nuclear Regulatory Research. Regulatory Guide 1.200.

Nuclear Regulatory Commission (1982). Crac-2 report nureg/cr-2239. Technical report.

Ottinger, R. L., Wooley, D. R., Robinson, N. A., Hodas, D. R., and Babb, S. E. (1990). *Environmental Costs of Electricity.* Oceana Publications Inc, Pace University Center for Environmental Studies.

Paté-Cornell, M. E. (1996). Uncertainties in risk analysis: Six levels of treatment. *Reliability Engineering & System Safety*, 54(2):95–111.

Rabl, A. and Rabl, V. A. (2013). External costs of nuclear: Greater or less than the alternatives? *Energy Policy*, pages 575–584.

Rasmussen, N. (1975). Wash-1400 (reactor safety study). *An assessment of accident risks in US Commercial Nuclear Power Plants*, 8.

Savage, L. J. (1954). *The Foundations of Statistics.* Dover Publications Inc., New York.

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

Slovic, P., Fischhoff, B., and Lichtenstein, S. (1982). *Understanding perceived risks*, pages 462–492. Cambridge University Press, Cambridge.

Slovic, P., Fischhoff, B., Lichtenstein, S., Corrigan, B., and Combs, B. (1977). Preference for insuring against probable small losses: Insurance implications. *The Journal of Risk and Insurance*, 44:237–258.

Sovacool, B. K. (2008). The costs of failure: A preliminary assessment of major energy accidents,1907-2007. *Energy Policy*, pages 1802–1820.

Van den Hout, W. B. (2010). The value of productivity: human-capital versus friction-cost method. *Annals of of the rheumatic diseases*, pages i89–i91.

Versicherungsforen (2011). Calculating a risk-appropriate insurance premium to cover third-party liability risks that result from operation of nuclear power plants. Technical report, Leipzig.

Viscusi, W. K. (1993). The value of risks to life and health. *Journal of Economic Literature*, 31:1912–1946.

Von Neumann, L. J. and Morgenstern, O. (1947). *Theory of games and economic behavior*. Princeton University Press, Princeton, NJ.

Wang, Q., Chen, X., and Yi-Chong, X. (2013). Accident like the fukushima unlikely in a country with effective nuclear regulation: Litterature review and proposed guidelines. *Renewable and Sustainable Energy Reviews*, pages 126–146.

Weitzman, M. L. (2011). Fat-tailed uncertainty in the economics of catastrophic climate change. *Review of Environmental Economics and Policy*, 5(2):275–292.

WSP Canada Inc. (2013). Risk assessments for marine oil spills in Canadian waters. Technical report, Transport Canada. http://www.tc.gc.ca/eng/marinesafety/menu-4100.htm.