Tough enough? Robust satisficing as a decision norm for long-term policy analysis

Andreas Mogensen and David Thorstad

Global Priorities Institute | November 2020

GPI Working Paper No. 15-2020

GLOBAL PRIORITIES INSTITUTE

UNIVERSITY OF OXFORD

Tough enough? Robust satisficing as a decision norm for long-term policy analysis

Abstract: This paper aims to open a dialogue between philosophers working in decision theory and operations researchers and engineers whose research addresses the topic of decision making under deep uncertainty. Specifically, we assess the recommendation to follow a norm of robust satisficing when making decisions under deep uncertainty in the context of decision analyses that rely on the tools of Robust Decision Making developed by Robert Lempert and colleagues at RAND. We discuss decision-theoretic and voting-theoretic motivations for robust satisficing, then use these motivations to select among candidate formulations of the robust satisficing norm. We also discuss two challenges for robust satisficing: whether the norm might in fact derive its plausibility from an implicit appeal to probabilistic representations of uncertainty of the kind that deep uncertainty is supposed to preclude; and whether there is adequate justification for adopting a satisficing norm, as opposed to an optimizing norm that is sensitive to considerations of robustness.

1.

We as a species confront a range of profound challenges to our long-term survival and flourishing, including nuclear weapons, climate change, and existential risks from biotechnology and artificial intelligence (Bostrom and Cirkovic 2008; Ord 2020). Policy decisions made in answer to these threats will impact the development of human civilization hundreds or thousands of years into the future. In the extreme, they will determine whether humanity has any future at all.

Unfortunately, severe uncertainty clouds our efforts to forecast the impact of policy decisions over long-run timescales. The term 'deep uncertainty' has been adopted in operations research and engineering to denote decision problems in which our evidence is profoundly impoverished in this way (Marchau et al. 2019). The third edition of the *Encyclopedia of Operations Research and Management Science* defines 'deep uncertainty' as arising in situations "in which one is able to enumerate multiple plausible alternatives without being able to rank the alternatives in terms of perceived likelihood" or "what is known is only that we do not know." (Walker, Lempert, and Kwakkel 2013: 397) A multi-disciplinary professional organization dedicated to research on decisionmaking under deep uncertainty (DMDU) was established in December 2015.

Given the paucity of evidence that could constrain probability assignments in these contexts, analysts associated with the study of DMDU argue that orthodox approaches to decision analysis based on expected value maximization are unhelpful. Thus, Ben-Haim (2006: 11) suggests that "[d]espite the power of classical decision theories, in many areas such as engineering, economics, management, medicine and public policy, a need has arisen for a different format for decisions based on severely uncertain evidence." The DMDU community has developed a range of tools for decision support designed to address this need, including *Robust Decision Making* (RDM) (Lempert, Popper, and Bankes 2003), *Dynamic Adaptive Policy Pathways* (DAPP) (Haasnoot et al. 2013), and *Info-Gap Decision Theory* (IG) (Ben-Haim 2006).

In spite of what a name like 'Info-Gap Decision Theory' might suggest, these tools are not faithfully characterized as decision theories, at least not in the sense in which philosophers would most naturally understand that term. They primarily comprise procedures for framing and exploring decision-problems. Sometimes their proponents seem reluctant to suggest normative criteria for solving decision-problems once suitably framed. Thus, Lempert et al. (2006: 523) write: "RDM does not determine the best strategy. Rather, it uses information in computer simulations to reduce complicated, multidimensional deeply uncertain problems to a small number of key trade-offs for decision makers to ponder."

This is the aspect of DMDU research highlighted in a recent paper by Helgeson (2020), one of the few papers by philosophers to examine research in this area in depth. Whereas decision theory, as practiced by philosophers, characteristically aims to instruct decision makers in how to solve a decision problem that is appropriately framed, Helgeson argues that DMDU research focuses principally on how to frame the decision problem in the first place, and so provides "a counterbalancing influence to decision theory's comparative focus on the choice task." (267)

DMDU researchers nonetheless also provide suggestions for normative criteria appropriate to the solution of the choice task. In particular, they tend to advocate a norm of *robust satisficing*. Lempert (2002: 7309) states that DMDU decision support tools "facilitate the assessment of alternative strategies with criteria such as robustness and satisficing rather than optimality. The former are particularly appropriate for situations of deep uncertainty." In the same vein, Schwartz, Ben-Haim, and Dasco (2011: 213) argue that "[t]here is a quite reasonable alternative to utility maximization. It is maximizing the robustness to uncertainty of a satisfactory outcome, or robust satisficing. Robust satisficing is particularly apt when probabilities are not known, or are known imprecisely." In a sense, it is unsurprising that conceptual innovations in decision framing should be accompanied by the proposal of novel decision criteria, as any reasonable decision support tool must make some assumptions about what constitutes a good decision in order to determine what information should be emphasized in framing the problem at hand.

Our interest in this paper is in robust satisficing as a norm for decision making under deep uncertainty. There has been remarkably little philosophical discussion of robust satisficing as a candidate decision norm, given its popularity among those at the coalface working on DMDU. Our aim in this paper is to open a dialogue between our two research communities. In order to simplify the discussion, we focus specifically on robust satisficing in the context of RDM and set aside other approaches, such as IG. The key questions on which we focus are how to characterize robust satisficing as a decision norm in the context of RDM, its relationship to more familiar decision criteria discussed among philosophers and economists, and whether in fact it is a rationally defensible norm that is suited to decision making under deep uncertainty.

2.

The two key goals of this section are to outline the procedures and decision support tools used to facilitate decision framing in the context of RDM, and to highlight and interpret key claims about robustness and satisficing as decision criteria made by researchers associated with its development. Along with the methodological remarks in section 3, this will set the stage for the subsequent critical discussion of robust satisficing in sections 4 and 5.

2.1

The procedures and decision support tools used in RDM build on the techniques of *scenario-based planning* (Schwartz 1996) and *assumption-based planning* (Dewar et al. 1993). Scenario-based planning encourages decision-makers to construct detailed plans in response to concrete, narrative projections of possible futures. Assumption-based planning emphasizes the identification of load-bearing assumptions, whose falsification would cause the organization's plan to fail.

Developed at the RAND corporation around the turn of the millennium, RDM uses computer-based exploratory modelling (Bankes 1993) to augment these techniques. Computational simulations are run many times, in order to build up a large database of possible futures and evaluate candidate strategies across the ensemble. This contrasts with the small handful of user-generated futures considered in traditional scenario planning, with Schwartz (1996: 29) recommending use of just two or three scenarios for the sake of tractability. Once an ensemble of futures has been generated, RDM uses scenario discovery algorithms to identify the key factors that determine whether a given strategy is able to satisfactorily meet the organization's goals. Data visualization techniques are used to represent the performance of alternative strategies across the ensemble of possible futures and identify key trade-offs.

By way of illustration, Robalino and Lempert (2000) apply RDM in exploring the conditions under which government subsidies for energy-efficient technologies can be an effective complement to carbon taxes as a tool for greenhouse gas abatement. The first step of their analysis is to construct a *system model*. Robalino and Lempert develop an agent-based model in which energy consumption decisions made by a heterogenous population of agents generate macroeconomic and policy impacts, which in turn affect the behaviour of agents at later times.

The second step is to settle on a set of *policy alternatives* to compare. Robalino and Lempert compare three strategies: a tax-only approach, a subsidy-only approach, and a combined policy using both taxes and technology subsidies.

Because model parameters are uncertain, the next step is to use data and theory to constrain model parameters within a plausible range. Sometimes this can be done with a point estimate: for example, the pre-industrial atmospheric concentration of carbon is estimated at 280 parts per million. Other parameters can be estimated only more imprecisely. For example, the elasticity of economic growth with respect to the cost of energy is constrained to the range [0, 0.1] (Dean and Hoeller 1992). The resulting *uncertainty space* of allowable combinations of parameter values is intractably large. As a result, a genetic search algorithm was used to identify a *landscape of plausible futures*: a set of 1,611 combinations of parameter values intended to serve as a good proxy for the larger uncertainty space.

To improve the tractability of analysis and the interpretability of results, a dimensionality-reduction approach was used to identify the five most decision-relevant model parameters. Allowing these parameters to vary across the landscape of plausible futures and fixing the rest at their mean values generates an *ensemble of scenarios* in which the performance of policy alternatives can be compared. In each scenario, alternatives were assessed by their *regret*: the difference in performance between the given alternative and the best performing alternative.

The purpose of RDM is to identify robust alternatives, which in this case means seeking alternatives with low regret across the ensemble of scenarios. This analysis leads to policy-relevant conclusions. For example, Robalino and Lempert find that a combined policy can be more effective than a tax-only policy if agents have heterogenous preferences, and there are either significant opportunities for energy technologies to deliver increasing returns for scale, or for agents to learn about candidate technologies by adopting them.

Although RDM models are not intended to deliver precise probabilistic conclusions, we can get a grip on the ways in which the likelihoods inform policy choice by treating the ensemble of scenarios as an equally-weighted collection of ways that the future could plausibly be. Proceeding in this way, Robalino and Lempert classify scenarios by the degree to which they exhibit high damages from climate change, and the degree to which the economy obeys classicality in assumptions about learning and returns to scale. They find that the combined strategy is preferable when both probabilities are moderate or large (Figure 1).

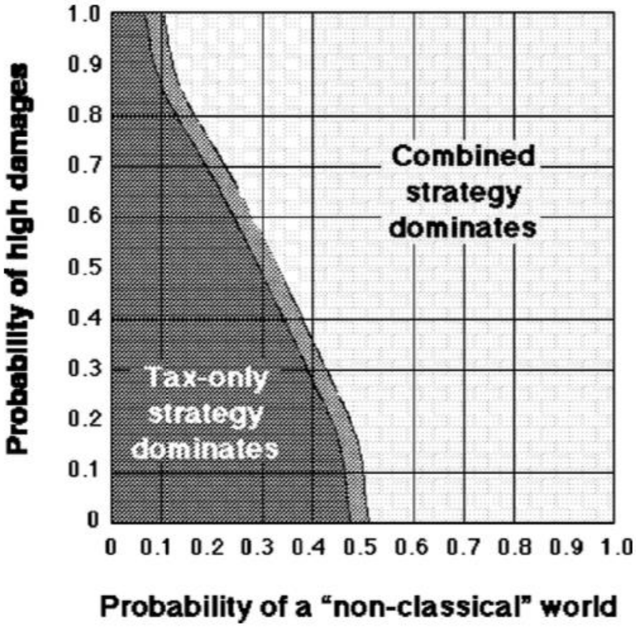

Figure 1: Performance of combined and tax-only strategies against climate damages and economic classicality – Reproduced from Robalino and Lempert (2000: 15)

Because RDM analysis uses a subset of the full uncertainty space, the last step is to search through the full uncertainty space for possible counterexamples to these strategy recommendations. If necessary, these counterexamples can be used to develop new alternatives or to constrain the landscape of plausible futures in a way more likely to discriminate between candidate alternatives, and then the analysis can be repeated.

2.2

The foregoing gives a sense of the procedures and tools used in decision framing that characterize the RDM approach. Our aim in this sub-section is to highlight and offer some preliminary interpretation of the recommendation to engage in robust satisficing when solving decision problems of the kind to which RDM is applied.

Researchers associated with the development of RDM typically emphasize the extreme fallibility of prediction in the face of deep uncertainty and recommend identifying strategies that perform reasonably well across the entire ensemble of possible futures, as opposed to fixing a unique probability distribution over scenarios relative to which the expected utility of options can be assessed. It is granted that expected utility maximization is appropriate when uncertainties are shallow. Lempert et al. (2006: 514) emphasize that "[w]hen risks are well defined, quantitative analysis should clearly aim to identify optimal strategies … based on Bayesian decision analysis … When uncertainties are deep, however, robustness may be preferable to optimality as a criterion for evaluating decisions." They define a robust strategy as one that "performs relatively well – compared to alternatives – across a wide range of plausible futures." (Ibid.)

In line with the example discussed in section 2.1, it is common for RDM theorists to use a *regret measure* to assess the performance of different strategies within a scenario. Regret may be measured in either absolute terms or in relative terms: i.e., as the difference between the performance of the optimal policy and the assessed policy or as this difference considered as a fraction of the maximum performance achievable. Using either measure, a robust strategy is defined by Lempert, Popper, and Bankes (2003: 56) as "one with small regret over a wide range of plausible futures, F, using different value measures, M." No details are given about how to determine what counts as 'small regret,' so far as we are aware. Decision makers who rely on RDM tools are presumably asked to define their own preferred satisficing threshold.

As the foregoing makes clear, the standard approach considers how well a given policy performs across the range of possible futures: i.e., its ability to produce satisfactory outcomes given different possible states of the world. So understood, robust satisficing could be considered as a competitor to standard non-probabilistic norms for decision making under what Luce and Raiffa (1957) term 'complete ignorance': i.e., as a competitor to norms like Maximin, Minimax Regret, or Hurwicz.

The assessment of a candidate policy against the norm of robust satisficing need not entirely forego probabilities, however. As we saw, Robalino and Lempert are willing to use probabilistic analysis to shed light on the ways in which economic non-classicality and damages from climate change bear on the optimality of taxes and subsidies. In this vein, Lempert, Popper, and Bankes (2003: 48) note that the notion of robust satisficing "can be generalized by considering the ensemble of plausible probability distributions that can be explored using the same techniques described here for ensembles of scenarios." In other words, we may choose to assess the robustness of a given policy by considering the extent to which its expected performance is acceptable across the range of probability distributions taken to be consistent with our evidence. So understood, robust satisficing may be considered as a competitor to norms for decision making with imprecise probabilities, like Liberal or MaxMinEU (discussed in section 4.3).

Notably, the norm of robust satisficing is presented by Lempert and Collins (2007) and Lempert (2019) as inspired by Starr's *domain criterion* for decisionmaking under complete ignorance, which canonically makes use of sets of probability assignments (Starr 1966; Schneller and Sphicas, 1983). Roughly speaking, the domain criterion asks agents who are completely ignorant about which world-state is actual to consider the set of all possible probability assignments to the states, giving equal weight to each possible probability assignment and choosing the act which is optimal on the largest number of probability assignments.[1]

Apart from the more explicit focus on deciding by reference to sets of probability values, we find the inspiration drawn from the domain criterion to be striking, in that it serves to highlight two concerns about robust satisficing that we will explore in greater depth in this paper.

The first concerns the relationship between robustness and satisficing. The domain criterion is recommended by Starr in part on the basis of considerations of robustness. Starr (1966: 75) argues that his criterion is superior to the Laplace criterion because, he claims, a decision rule that relies on a unique probability assignment to the possible states is too sensitive to the particular probability vector that we

[1] That is, the domain criterion adopts a uniform second-order probability measure μ over the set of coherent first-order probability functions and selects the option(s) which maximize utility on the μ-largest set of probability functions.

choose. Nonetheless, the domain criterion compares strategies in terms of the number of probability assignments on which they *maximize* expected utility and is therefore naturally thought of as a norm of robust *optimizing*, albeit not in the sense in which 'robust optimization' is standardly treated in textbooks on optimization theory (e.g., in Ben-Tal, Ghaoui, Nemirovski 2009). In expositions of RDM, robustness and optimizing are frequently contrasted, and a desire for robustness is linked to satisficing choice. Starr's domain criterion serves to highlight that there is no strictly logical connection here. What, then, explains the emphasis on satisficing choice that informs the design and application of RDM? We take up this issue in 5.

Secondly, we re-iterate that in Starr's exposition, the domain criterion explicitly relies on higher-order probabilities. We are "to measure the probability that a randomly drawn [first-order probability assignment] will yield a maximum expected value for each particular strategy." (Starr 1966: 73) This is done using a uniform second-order probability measure over first-order probability functions, a strategy Starr takes to represent a more plausible interpretation of the Principle of Insufficient Reason than the Laplace criterion. Notably, a similar approach is suggested by Rosenhead, Elton, and Gupta (1972), in what may be considered the *locus classicus* for discussions of robustness as a decision criterion in the field of operations research. In planning problems in which an initial decision d_i must be chosen from a set D and serves to restrict the set S of alternative plans capable of being realized in the long run to a subset S_i, they define the robustness of d_i in terms of $|\tilde{S}_i|/|\tilde{S}|$, where $|X|$ denotes the cardinality of set X and $\tilde{S} \subseteq S$ is the set of long-run outcomes considered "acceptable according to some combination of satisficing criteria." (419) Rosenhead, Elton, and Gupta explicitly note that using their robustness criterion is equivalent to using the Laplace criterion for decision making under complete ignorance if the agent's utility function is defined to be 1 for any $s \in \tilde{S}$ and 0 otherwise.

Given its pedigree, it would seem natural to expect that the robustness criterion appealed to in RDM must also involve reliance on a uniform first-order or second-order probability measure, if only implicitly. However, this seems at odds with the skepticism toward unique probability assignments that is otherwise characteristic of the DMDU community (see, e.g., Lempert, Popper, and Bankes 2003: 47-8). We consider this issue in greater depth in Sections 4.4-4.5.

As one final point, we note that many formulations of the robust satisficing norm emphasize robustness not only with respect to different possible states of the world or different probability assignments, but also with respect to alternative valuations of outcomes. Thus, Lempert, Popper, and Bankes (2003: 56) state that a robust strategy must perform reasonably well as assessed "using different value measures". For the sake of simplicity our discussion throughout this paper will ignore the issue of robustness with respect to different value systems and focus exclusively on robustness with respect to empirical uncertainty.

3.

Before we proceed to the critical evaluation of robust satisficing as normative principle for decisionmaking, we pause briefly to reflect on an important methodological issue, which relates to the distinction between a *criterion of rightness* and a *decision procedure*, familiar to moral philosophers from debates on utilitarianism (see, *inter alia*, Bales 1971, Hare 1981, Railton 1984).

As moral philosophers understood the term, a criterion of rightness specifies, in the broadest possible terms, the conditions under which a given act is right or wrong. It is natural to view normative standards of this kind as necessary and *a priori*. Adoption of a given criterion of rightness need not commit us to thinking that people in general should deliberate by considering whether the conditions specified by the criterion are satisfied. This need not be a reliable means of choosing those acts that satisfy the criterion. Whether a given decision procedure is in fact a reliable decisionmaking tool for a given agent in a given context depends on factors that are contingent and *a posteriori*.

Orthodox normative decision theory is arguably in the business of specifying *a priori* criteria of rational choice. By contrast, research on fast and frugal heuristics (Gigerenzer, Todd, and the ABC Research Group 1999) is explicitly concerned with what decision procedures work best in practice when deployed by boundedly rational agents, emphasizing the extent to which practical success is contingent on local features of the environment in which a heuristic is applied.

How should we understand the recommendation to engage in robust satisficing when making decisions under deep uncertainty? It may seem natural to some to think that this must be a decision heuristic whose reliability is contingent and *a posteriori*. This interpretation may be bolstered by the assertion by Lempert et al. (2006: 518) that, "[u]ltimately, the claim that RDM will help decision makers make better decisions in some important situations than traditional approaches must be empirically tested."

If RDM or robust satisficing are treated as contingently reliable heuristics, then they should be supported by *a posteriori* evidence of successful decisionmaking. However, the methods are too young to have a suitable track record, being recently developed tools for making policy decisions whose impacts will play out over long-run timescales. Hence, if RDM and robust satisficing are treated as contingently reliable heuristics, we would have little evidence to justify their use. It could be protested that there is ample empirical evidence of the failure of traditional decisionmaking methods under conditions of deep uncertainty (Goodwin and Wright 2010), but this evidence would not discriminate finely enough to favour RDM or robust satisficing over other extant or yet-to-be-developed methods.

Insofar as robust satisficing is appealing as a decision norm, we think its appeal must reflect *a priori* intuitions about the criteria for choice under deep uncertainty. Since we don't know whether it works in practice, all we have to go on is whether it works in theory. The central question on which we focus in the remainder of this paper are therefore to do with the nature and credibility of the intuitive *a priori* considerations that may be taken to support robust satisficing as a decision norm.

4.

In this section, we explore what we take to be the core consideration put forward in the RDM literature for preferring decision norms that exhibit robustness, in the sense(s) outlined in section 2. This takes the form of an intuitive objection to the application of subjective expected utility theory to problems involving deep uncertainty. This objection receives varied expression in the RDM literature. We show that it generalizes to other decision norms besides subjective expected utility maximization, but also serves as a flashpoint for the worry, mentioned in section 2.2., that robust satisficing invokes a form of second-order probability theory.

4.1

Consider the approach suggested by Broome (2012) for policy making in the face of global climate change. Broome concedes that our evidence about climate impacts does not uniquely constrain the probabilities to be assigned to all relevant contingencies. His recommendation is that we should nonetheless assign these contingencies precise probabilities on the basis of a subjective best-guess estimate and maximize expected value relative to that assignment. Broome writes: "Stick with expected value theory, since it is very well founded, and do your best with probabilities and values." (129)

The core worry that we perceive as animating the emphasis on robustness in the RDM literature is that a policy that is optimal relative to a subjective best-guess probability assignment may in principle be suboptimal, or even catastrophic in expectation, relative to the other probability assignments that were not excluded by our evidence. Intuitively, when this is the case, it is unwise to optimize relative to one's subjective best-guess probability estimate (compare Sprenger 2012), and decisionmaking may be improved by incorporating alternative or less committal ways of representing uncertainty. This view is often adopted in discussions of climate change. For example, the most recent report of the Intergovernmental Panel on Climate Change (IPCC) communicates uncertainty using five qualitative levels of confidence, ranging from `very low' to `very high', together with a ten-point likelihood scale, with each term calibrated to correspond to an interval of precise probabilities (IPCC 2010, 2014). This reflects the IPCC's view that, given the depth of climate uncertainties, decisionmakers may benefit from considering a range of plausible futures or likelihoods rather than optimizing policies to match a deeply uncertain best-guess probability assignment (Bradley et al 2017; Helgeson et al 2018).

4.2

On its face, the intuitive objection raised against subjective expected utility maximization in 4.1 is also an objection to many standard norms for decision making that reject precise probabilism and assume instead that the decision maker's doxastic state is modelled by a (non-singleton) set of probability functions, R, often called a *representor* (following van Fraassen 1990). The representor is sometimes imagined metaphorically as a 'credal committee' (e.g., in Joyce 2010): an assembly of agents each of whom has a definite degree of confidence for each proposition, corresponding to each probability function in R. Within this framework, deep uncertainty is characterized by a wide representor containing many different probability functions.

Consider the Liberal norm, which states that an act is rationally permissible just in case it maximizes expected utility relative to some $\Pr(\cdot) \in R$. Echoing the points noted in 4.1, it feels natural to object to the Liberal rule on the basis that it may be unwise to choose an act that is optimal relative to some arbitrary probability distribution over the possible states, without regard for whether the act may be evaluated very differently across the other probability assignments consistent with one's evidence. For example, in an imprecise probabilistic recasting of the study by Robalino and Lempert, the Liberal norm would likely say that any of the candidate policies – taxation, subsidies, or a mixture of the two – is permissible, because for each policy there is some probability function consistent with the evidence that recommends this policy.

The same concern extends to the Maximal rule, which states that an act is rationally permissible just in case there is no other act with greater expected utility relative to every $\Pr(\cdot) \in R$, since any act that is permissible according to Liberal is permissible according to Maximal. These decision rules seem to give individual probability functions too much power by allowing them to veto any recommendation against choice of a given policy that would otherwise be made by one's 'credal committee'. As before, the Maximal rule would probably imply that any of the candidate policies for technology diffusion is permissible because no policy outperforms any other across all admissible probability assignments.

A similar concern arguably extends to many other decision rules commonly discussed in the literature on imprecise credences, such as MaxMinEU and HurwiczEU. MaxMinEU ranks acts in terms of their minimum expected utility relative to the probability functions in R. HurwiczEU is a generalisation of MaxMinEU that ranks the available acts in terms of a convex combination of their minimum and maximum expected utilities relative to the probability functions in R. Intuitively, both rules give extremal probability functions too much power in settling the recommendation of one's 'credal committee'. Intuitively, we would like to object that policies should also be evaluated by their performance on more moderate probabilistic hypotheses, which give middling credence to market failures and other forms of non-classicality.

The charge made in the previous paragraph was that HurwiczEU ignores moderate probability functions, letting decisionmaking be driven only by the pair of most pessimistic and optimistic probability functions, respectively. One would like to put the point as follows: HurwiczEU does not change its recommendations about what should be done if all of the non-extremal probability functions are removed from the agent's representor. But this way of putting the point is too hasty, since it is often held, following Levi (1980), that representors should be convex. Hence it does no good to object that HurwiczEU would not change its recommendations if these moderate mixtures were removed, because if convexity is a rational requirement, then it would be irrational for a decisionmaker to remove them. But we can put the point in another way. Suppose that o and o' are tax policies which perform comparably well in the classical limit and comparably poorly as market failures accumulate. But suppose that o, unlike o', performs well under moderately nonclassical conditions. Intuitively, we would like to say that o is preferable to o', because o is more robust to economic uncertainty. But HurwiczEU ranks o and o' comparably, and if o' outperforms o even slightly under extremal conditions, HurwiczEU may recommend o' over o. This suggests that HurwiczEU, as well as the MaxMinEU rule that it generalizes, may place insufficient emphasis on robustness to deep uncertainty.

4.3

In formulating this line of objection as it applied to subjective expected utility maximization or the Liberal rule for imprecise credences, we emphasized that a policy that is optimal relative to some particular admissible probability assignment may in principle be suboptimal or even catastrophic in expectation relative to other probability assignments that are not excluded by our evidence. Intuitively, the objection is not simply that there is *some* other admissible probability function relative to which the policy fails to maximize expected utility. Thus, the intuitions to which we have appealed are plausibly not of the kind that would lead us to accept the so-called Conservative rule, according to which an act is rationally permissible just in case it maximizes expected utility relative to *every* $\Pr(\cdot) \in R$. Intuitively, the Conservative rule evokes the same concern about giving veto power to individual members of the agent's 'credal committee' as the Liberal rule, albeit in the opposite direction. In our case study based on Robalino and Lempert (2000), the Conservative rule would presumably imply that no government policy is permissible because for each policy, some member of the representor recommends against that policy.

Plausibly, the intuitive objection that we have sought to articulate in 4.1-4.2 gains purchase only insofar as it makes sense to say things such as that there are *more* admissible probability functions on which the policy fails to maximize expected utility than on which it succeeds. The objection to rules such as Liberal, Maximal, MaximinEU and HurwiczEU is that they heed the advice of two or fewer probability functions, even when many more elements of the agent's `credal committee' disagree. However, one might worry that this talk of *more* and *fewer* admissible probability functions cannot be made precise without helping ourselves to a second-order probability measure on the set of admissible probability assignments, by which to quantify the size of a collection of first-order probability distributions.

For reasons we've already noted, this would, in a sense, be unsurprising. The robust satisficing norm is characterized in the RDM literature as a descendant of Starr's domain criterion, which explicitly relies on a uniform second-order probability function. In a different sense, however, reliance on a (unique) second-order probability function would be very surprising for anyone driven by the sort of concerns that motivate RDM and the DMDU community as a whole. Prominent among them is the thought that it is inappropriate to require decision makers facing severely uncertain long-term policy problems to express their uncertainty in the form of a unique probability distribution. Lempert (2002: 7310) notes that the decision support tools associated with RDM are designed "to help users form and then examine hypotheses about the best actions to take in those situations poorly described by well known probability distributions and models."

These remarks are to do with first-order probabilities. Nonetheless, it seems implausible to require precise second-order probabilities while conceding that our evidence is so incomplete and/or ambiguous that no precise probability assignment over the relevant first-order hypotheses is warranted by our evidence. In fact, there is a strong case to be made that this is simply incoherent.

We have in mind here the familiar objection, pressed by Savage (1972: 58), that probabilistic uncertainty with respect to the correct probability function over first-order hypotheses requires assigning a unique probability to each first-order hypothesis, corresponding to the (second-order) expectation of the first-order probability.[2] Thus, suppose that the second-order probability represents the agent's subjective probability with respect to the uniquely rational credence assignment warranted by her evidence. It seems very plausible that any rational agent's subjective probabilities should defer to the uniquely rational credence function, so understood, in the sense that if $\mathrm{Pr}(\cdot)$ is the subjective probability function for an agent with total evidence E and $\mathrm{Ev}_E(\cdot)$ is the uniquely correct credence

[2] Responding to Savage, Skyrms (1980: 117) and Sahlin (1993: 25) argue that if we attach significance not simply to the expectation, but also to the overall shape of the higher-order probability distribution (e.g., to its variance or skew), then higher-order probabilities need not be otiose. However, we are not arguing that higher-order probabilities are superfluous, but simply that plausible deference principles governing unique assignments of higher-order probabilities require assigning unique subjective probabilities to first-order propositions.

function warranted by E, then we require that $\Pr(H|\ \mathrm{Ev}_E(\ H) = x) = x$ for any hypothesis H.[3] It then follows by the Law of Total Probability that $\Pr(H) = \mathbb{E}[\mathrm{Ev}(H)]$, where the expectation is defined relative to $\Pr\,(\,\cdot\,)$. Thus, the agent, if she is rational, has some precise degree of confidence in H after all. Similar deference principles concerning the chance function (Lewis 1980) or the agent's own credences (van Fraassenn 1984) yield similar results if we assume that the relevant higher-order probabilities represent the agent's credences for different chance hypotheses or the agent's credences with respect to what her own credences are.

4.4

In light of the foregoing, we think proponents of robust satisficing have some explaining to do about the relationship between their position and second-order probability theory. If the exposition and defence of robust satisficing relies on second-order probabilism, then we may find ourselves committed to a unique first-order probability assignment after all.

This problem may be answerable. In particular, there is the following rejoinder to the objection that a decision criterion formulated in terms of a robustness measure, μ, on R presupposes second-order probabilities: namely, that μ cannot be a probability measure because R is not a *sample space*. Roughly speaking, a sample space is a set of different ways things could be, exactly one of which is realized. However, R need not be viewed as a set of probability functions exactly one of which is true or correct.

It could be suggested that there is always some credence function that is uniquely warranted by our evidence, whereas you and I lack the cognitive wherewithal to fix our credences with the required degree of precision. The elements of the set of probability functions that represents our

[3] To say that this principle is very plausible is not to suggest that it is without critics. Dorst (2019) notes that for agents who know their actual credences, this deference principle is incompatible with a certain form of higher-order uncertainty, which Dorst finds unacceptable.

doxastic state might then naturally be interpreted as different candidates for the one true rational credence function.

However, this sort of view is atypical in the recent philosophical literature on imprecise credences. Proponents of imprecise probabilism like Joyce (2005, 2010) instead take the view that even an ideally rational agent will, in some cases, refrain from assigning precise probabilities. Thus, Joyce (2010: 283) insists that "since the data we receive is often incomplete, imprecise, or equivocal, the epistemically *right* response is often to have opinions that are similarly incomplete, imprecise or equivocal." In particular, it is not the case that some element of R represents the epistemically right response. That role is to be played by the mushy credal state represented by R itself.

If we take a view of this kind, then R cannot be interpreted as a sample space, since its elements are not different candidates for the uniquely correct probability function. As a result, a measure on R is not a probability measure. In this way, one might argue, a demand for robustness does not presuppose second-order probabilities.

We are not sure how persuasive this rejoinder will be found to be. Some may worry that, for all we have said, the intuitive appeal of relying on a measure on R derives from implicitly treating R as a sample space and the measure as a probability measure. We also note the following. It is relatively common in the RDM literature – indeed, throughout the DMDU literature - to describe cases in which decision makers' uncertainty about the future is not appropriately represented by a unique probability distribution as cases in which probabilities are *unknown*. Lempert, Popper, and Bankes (2003) write that since "traditional decision theory addresses the multiplicity of plausible futures by assigning a likelihood to each and averaging the consequences" (26), traditional decision theory is unhelpful for long-term policy analysis because there is "simply no way of knowing the likelihood of the key events that may shape the long-term future." (27) Similarly, Marchau et al. (2019: 7) describe scenario-based planning as appropriate for cases in which "the likelihood of the future worlds is unknown". These remarks are naturally taken to imply that there is some uniquely correct probability function among the set of probability functions that we take our evidence as failing to rule out. It is just that we don't know what it is. In that case, we cannot argue, in the way we have done here, that the set cannot be construed as a sample space.

Here is a second point, in a similar vein. Readers may have noted that our discussion throughout this section has focused on robustness considered with respect to an ensemble of probability distributions, as opposed to the more common emphasis in RDM analyses on robustness with respect to the range of plausible future outcomes. Understood in this way, robust satisficing is to be considered as akin to non-probabilistic norms for decision making under complete ignorance, like Maximin, Minimax Regret, or Hurwicz, and not to decision norms for imprecise probabilities like MaxMinEU or HurwiczEU.

As noted in our discussion of Rosenhead, Elton, and Gupta (1972) on robustness as a decision criterion, robust satisficing, so understood, also seems to require a measure over the set of states that behaves like a (uniform) probability distribution. Furthermore, in this context, there is no plausibility to the idea that the set which forms part of the measure space cannot be interpreted as a sample space. Clearly, exactly one of these states is the true state of the world. The objection that we are here just smuggling in probabilities via the backdoor strikes us as harder to evade.

5.

Our discussion in section 4 concerned the desirability of choosing options that are robustly supported, in the sense of being approved by a larger proportion of the probability functions in R. In this section, we explore the relationship between robustness and satisficing. As noted, in expositions of RDM, robustness and optimizing are frequently contrasted, and a desire for robustness is linked to satisficing choice. However, there is no straightforward relationship between these concepts, and Starr's domain criterion is naturally characterized as a criterion for robust optimizing.

5.1

Recommendations to adopt a satisficing norm are sometimes likened in the RDM literature to the notion of satisficing choice discussed by Simon (1955), e.g., by Lempert, Popper, and Bankes (2003: 54) and Lempert et al. (2006: 516). This may suggest that robust satisficing is recommended for the same reasons that recommend a satisficing approach in Simon's work on bounded rationality. However, we find this implausible.

It is important to distinguish satisficing in a static setting from satisficing in a dynamic setting (Schmidtz 2004). In Simon's treatment, satisficing is a search heuristic for a class of dynamic decision problems involving serially presented options. It instructs decisionmakers to fix an aspiration level and halt their search when they find an option that meets that level, thereby evading the costs associated with added search and/or iteratively computing the expected utility of further search (Selten 1998).

The application of a satisficing stopping rule in search problems is to be contrasted with satisficing choice in static contexts, where the available options are open to view all at once. The rationality and/or moral permissibility of satisficing choice in static contexts is what was at issue in the debate between Slote (1984, 1989) and his critics (see Byron 2004). Slote argued that agents are sometimes rationally and/or morally permitted to choose a satisfactory outcome even when a better outcome is known to be available. For example, someone in a fairy-tale who is given the opportunity to make a wish on his death bed is supposed to be able to rationally wish for his family to be comfortably well-off forever after, even though he could have wished for them to be better off and would have preferred that they were.

RDM is often applied in contexts that do not have a fixed menu of candidate strategies. Lempert, Popper, and Bankes (2003: 66) note that "human participants in the analysis are encouraged to hypothesize about strategic options that might prove more robust than the current alternatives … These new candidates can be added to the scenario generator and their implications dispassionately explored by the computer." Thus, there is a search component in practical applications of RDM. Nonetheless, the robust satisficing decision rule recommended by Lempert and colleagues functions, so far as we can tell, as a criterion for the synchronous comparison of different candidate strategies. Asked to assess two candidate policies side-by-side in an imprecise probabilistic setting, we are to prefer that which does better in terms of performing satisfactorily across a wider range of admissible probability distributions over the possible futures. This recommendation is to be contrasted with Starr's recommendation to prefer the act that does better in terms of performing optimally across a wider range of admissible probability distributions.

This suggests that robust satisficing is more in line with Slote's conception of satisficing choice. This could lead readers to worry that robust satisficing has similarly contentious implications. Slote's conception of rational decision making permits the choice of dominated options. If we interpret the norm of robust satisficing as permitting choice of any strategy that performs sufficiently well against a wide range of admissible scenarios or probability distributions, then robust satisficing would also seem to permit the choice of dominated options. It would perhaps be permissible to choose any climate abatement policy that would keep emissions and expenses within satisfactory levels in a sufficiently high number of plausible futures, even over another policy that would reduce emissions or expenses further in every plausible future.

While we think that this worry lends urgency to the need for a precise formulation of the norm of robust satisficing, we do not think that the norm must be understood in a way that permits choosing dominated options. Most straightforwardly, we could hold that robust satisficing is only a necessary criterion on rational choice, allowing the choice between multiple robustly satisfactory options to be made on other grounds, such as ruling out dominated options.

5.2

A different way to justify the emphasis on satisficing in the RDM literature that we now wish to explore is by appeal to the study of voting methods.

In the literature on decision making and imprecise probabilities, comparisons with social choice are abundant. For example, Weatherson (1998: 6) writes: "we can regard each of the [probability functions in the representor] as a voter which voices an opinion about which choice is best". Viewed in this perspective, it is natural to consider voting rules as ways of aggregating preferences across members of the 'credal committee'. The major impediment to doing so presumably comes in the form of doubts about the use of a measure that would allow us to compare the size of different subsections of the 'electorate', which we addressed in the previous section.

We think the domain criterion is naturally interpreted as a form of *plurality voting*. Plurality voting is the most commonly used voting method in democracies today. Each voter casts a ballot for her most preferred option, and the option with the most votes is the winner. We shall interpret the domain rule as stating that when the agent's beliefs are modelled by a representor R and μ represents an indifferent weighing function over R, then any option o is valued at $\mu(M_o)$, the number of probability functions in the set M_o on which o maximizes expected utility, and choice of o is permissible just in case it has highest value, or is tied for highest value. Under the domain criterion, each admissible probability assignment may therefore be conceived as a voter who casts a ballot for her most preferred option, understood as that option that maximizes expected utility relative to the probability function whose vote is to be cast. So understood, the domain criterion requires choice of the option(s) with the greatest share of first-place votes.

An alternative to plurality voting is *approval voting* (Brams and Fishburn 1983). Ballots no longer indicate each voter's most preferred option, but rather all options that she finds acceptable, allowing each elector to vote for multiple candidates. Every ballot on which an option appears gives that option one vote. The election goes to the candidate(s) with the most votes. We find it natural to interpret the norm of robust satisficing suggested in the RDM literature as a kind of approval voting. In the imprecise probabilistic context, we interpret this decision rule as valuing options at $\mu(S_o)$, the number of probability functions in the set S_o on which o has satisfactory expected utility, with choice of a given option permitted just in case it has highest value, or is tied for highest value. In this case, each admissible probability assignment may be conceived as a voter who casts a ballot for every option that she finds satisfactory given the probability assignment whose vote is to be cast, with the winner as the option(s) with the most approval votes.

Earlier, we asked why the RDM literature links satisficing and robustness, as opposed to directing us to pick the option that is robustly optimal à la the domain rule. One way in which this question may be understood, we now suggest, is in terms of a choice between plurality voting and approval voting as a means of aggregating evaluations across the agent's 'credal committee'. The literature on voting methods contains a number of arguments for the relative superiority of approval voting (Brams and Fishburn 1983; Brams 2008). While some of these arguments fail to carry over to the current setting, others can be transposed, given suitable tweaks. We discuss one such argument below.

5.3

Ideally, a voting rule should avoid problematic forms of vote-splitting. Suppose that a diverse electorate is asked to choose between taxation and a combined tax/subsidy hybrid policy for climate abatement. 60% of the electorate favours the hybrid, and only 40% of the electorate favours the tax. Under plurality voting, the hybrid will be chosen. But now suppose that voters are asked to choose between a tax-only policy and a pair of hybrids, one of which includes a more generous subsidy than the other. 40% of voters continue to favour the tax, but the remaining 60% of voters split evenly between the two hybrid options. Now, the tax-only approach wins under plurality voting.

That seems like the wrong result. The tax-only approach has been chosen only because votes against the tax policy were split between a pair of competing policies. Fortunately, approval voting allows us to avoid this result. It does not require voters to choose between the two hybrid options. Most voters who approve of one mixed policy would presumably approve of the other, and vote-splitting would thus be avoided. This provides a strong reason to prefer approval voting to plurality voting.

Cashing out the thought that approval voting is less vulnerable than plurality voting to perverse forms of vote-splitting places important constraints on how approval voting and vote-splitting are understood. On the one hand, approval voting must not be understood so broadly that it collapses into plurality voting. This would happen, for example, if we allowed voters to approve of only their most preferred candidate. On the other hand, immunity to vote-splitting must not be understood so narrowly that it is satisfied even by this latter form of approval voting. This rules out some weak conceptions of immunity to vote-splitting, such as the revised conception of *clone independence* put forward by Tideman (2006: 174).

One natural type of immunity to vote-splitting is captured by Sen's condition α. Where S is the set of electoral candidates, P is a set of preference orderings, and $f(P, S)$ is a voting function that returns a subset of S as the winner(s) given P, f satisfies α just in case whenever $x \in f(P, S)$ and $x \in S' \subset S$, it is the case that $x \in f(P, S')$. In other words, a winning option would not have lost had we eliminated a different option from the ballot. Equivalently, no losing option would have won had some new alternative been added to the ballot. Thus, when f satisfies α, electoral outcomes are immune to the *spoiler effect*, whereby the entry of a new candidate into a political contest alters the electoral outcome in a way that favours an existing candidate rather than the entrant. As the example from the start of this section illustrates, plurality voting violates this condition.

Approval voting satisfies α so long as what counts as an approvable choice does not vary with the menu of options: for example, if any voter votes for any policy option just in case it yields (expected) GDP per capita above some fixed absolute threshold. Within the context of choice under deep uncertainty, the same can be said of a robust satisficing rule that scores options in terms of the value of $\mu(S_o)$ provided that $\Pr \in \mu(S_o)$ just in case o's expected value along some dimension (or composite of dimensions) relative to Pr exceeds some fixed absolute threshold. Call this *fixed threshold robust satisficing*.

In our view, vote splitting is no less of a concern in the context of decision making under deep uncertainty, with voters conceived as members of the decision maker's 'credal committee'. If anything, it is a greater concern. 'Voters' and 'candidates' in this context cannot make strategic choices to mitigate spoiler effects. Furthermore, the individuation of options is to a large extent arbitrary. A given policy can be subdivided indefinitely into its possible implementations, replacing one candidate with a multitude of clones. Insensitivity to spoiler effects in this context can be conceived as another kind of robustness: namely, as robustness with respect to arbitrary choices about how to individuate the decision maker's options. We take this to provide a strong reason to favour fixed threshold robust satisficing over the domain rule.

As we have already noted, the RDM literature does *not* typically define satisficing in terms of satisfaction of an absolute threshold. Instead, satisficing is typically understood in terms of a regret measure. As a result, what counts as a satisfactory choice can in principle vary with the menu of options, and the addition or subtraction of candidates from the menu of options can, respectively, change losers to winners or winners to losers. We think this provides a reason to reject the standard conception of the satisficing threshold appealed to in the RDM literature.

Of course, it may be objected that it is undesirable in some other respect to insist that decision makers should decide in light of a fixed absolute threshold. Lempert, Popper, and Bankes (2003: 56) suggest that a regret measure is to be recommended because this "supports the desire of today's decisionmakers to choose strategies that, in retrospect, will not appear foolish compared to available alternatives." For our part, we think insensitivity to spoiler effects should weigh much more heavily. In addition, we think that there are good independent reasons to operate with an absolute satisficing criterion when considering issues in long-term policy analysis related to global catastrophic risks like those enumerated in the introduction, where the avoidance of catastrophe is paramount. A catastrophe, in the sense at issue here, is arguably not merely some outcome that is sufficiently worse than the best that might have been achieved. Thus, we would not hesitate to describe the unanticipated occurrence of a runaway greenhouse effect that renders our planet inhospitable to life as catastrophic just in case we simultaneously discovered that this outcome was by now unavoidable. It seems more plausible to understand a catastrophe, in the sense at issue here, as an event that is bad relative to some threshold that does not depend on what actions are available.

5.4

The previous sub-section shows how the analogy with voting theory can be used not only to motivate robust satisficing over competing choice rules, but also to decide between candidate formulations of robust satisficing. Properties like satisfaction of Sen's α that adequately formalize the idea of insensitivity to spoiler effects will inevitably be shared by some, but not all versions of robust satisficing, and by tracing these properties across competing formulations we can begin to discriminate between them, albeit in a way that may seems to tell against the standard approach to defining the satisficing threshold adopted in the RDM literature.

We emphasize that the previous section offers just a taste of what we think voting theory can offer the study of robust decisionmaking in the face of deep uncertainty. There is more that could be said in comparing plurality voting and approval voting in this context. For example, another advantage claimed for approval voting over plurality voting by Brams and Fishburn (1983) is its greater propensity to elect Condorcet winners: i.e., options that are preferred by a majority to every other option in any pairwise comparison. In addition, there is no reason in principle to restrict our analysis to just these voting methods. This, we think, is ultimately the strongest obstacle to justifying an approach that marries a concern for robustness with an emphasis on satisficing, when it comes to long-term policy decisions involving imprecise probabilities. Even if approval voting in this context is preferable to plurality voting, it may well not be preferable to other voting methods that, like plurality voting, do not rely on any satisficing threshold to define when a candidate receives a vote.

6.

Our aim in this paper has been to open a dialogue between philosophers and the DMDU community about robust satisficing, specifically when applied as part of RDM decision analysis. We think philosophers can learn interesting lessons by considering the approaches suggested by the DMDU community for solving the choice task, and not merely its contributions to the art of decision framing.

We think philosophers have something to offer in return, and we hope this paper helps to make our case.

As befits an attempt to open a dialogue, we do not draw any hard and fast conclusions. We've noted objections that may be raised against robust satisficing as a decision criterion and suggested how these objections could be met, albeit in ways that require revising aspects of the standard presentation in the RDM literature. As we see it, the overarching upshot of our discussion is that the robust satisficing norm repays further clarification and assessment.

Bibliography

Bales, R. Eugene (1971), "Act-Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?," *American Philosophical Quarterly* 8.3: 257-265.

Bankes, Steve (1993), "Exploratory Modeling for Policy Analysis," *Operations Research* 41.3: 435-449.

Bradley, Richard and Helgeson, Casey and Brian Hill (2017), ``Climate Change Assessments: Confidence, Probability and Decision'', *Philosophy of Science* 84.3: 500-522.

Brams, Steven (2008) *Mathematics and Democracy: Designing Better Voting and Fair-Division Procedures*. Princeton, NJ: Princeton University Press.

Brams, Steven and Fishburn, Peter (1983) *Approval Voting*. Boston, MA: Birkhauser Press.

Ben-Haim, Yakov (2006), *Information Gap Decision Theory: Decisions under Severe Uncertainty,* 2nd edition. Cambridge, MA: Academic Press.

Ben-Tal, Aharon and El Ghaoui, Lauren and Arkadi Nemirovski (2009) *Robust Optimization*. Princeton, NJ: Princeton University Press.

Bostrom, Nick and Cirkovic, Milan (2008), *Global Catastrophic Risks*. Oxford, UK: Oxford University Press.

Broome, John (2012), *Climate Matters: Ethics in a Warming World*. New York, NY: WW Norton & Co.

Byron, Michael (2004), *Satisficing and Maximizing: Moral Theorists on Practical Reason*. Cambridge, UK: Cambridge University Press.

Dean, Andrew and Hoeller, Peter (1992), "Costs of Reducing CO2 Emissions: Evidence from Six Global Models," Working Papers, Department of Economics, no. 122, OECD, Paris.

Dewar, James, Builder, Carl, Hix, William and Levin, Morlie (1993), *Assumption-Based Planning: A Planning Tool for Very Uncertain Times*. Santa Monica, CA: RAND Corporation.

Dorst, Kevin (2019) Higher-Order Uncertainty. In Skipper and Steglich-Petersen, eds. *Higher-Order Evidence: New Essays*, 35-61. Oxford, UK: Oxford University Press.

Gigerenzer, Gerd, Todd, Peter, and the ABC Research Group (1999), *Simple Heuristics That Make Us Smart*. Oxford, UK: Oxford University Press.

Goodwin, Paul and George Wright (2010), "The Limits of Forecasting Methods in Anticipating Rare Events," *Technological Forecasting and Social Change* 77.3: 355-368.

Haasnoot, Marjolijn and Kwakkel, Jan H. and Walker, Warren E. and ter Maat, Judith (2013), ``Dynamic Adaptive Policy Pathways: A Method for Crafting Robust Decisions for a Deeply Uncertain World,'' *Global Environmental Change* 23.2: 485-498.

Hare, Richard (1981), *Moral Thinking: Its Levels, Method, and Point*. Oxford, UK: Clarendon Press.

Helgeson, Casey (2020): ``Structuring Decisions under Deep Uncertainty,'' *Topoi* 39: 257-269.

Helgeson, Casey and Bradley, Richard and Brian Hill (2018), "Combining Probability with Qualitative Degree-of-Certainty Metrics in Assessment,'' *Climactic Change* 149.3-4: 517-525.

IPCC (2010). *Guidance Note for Lead Authors of the IPCC Fifth Assessment Report on Consistent Treatment of Uncertainties*. Available at <http://www.ipcc.ch>.

IPCC (2014), *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.* Geneva, Switzerland: IPCC.

Joyce, James M. (2010) "A Defense of Imprecise Credences in Inference and Decision Making," *Philosophical Perspectives* 24.1: 281-323.

Joyce, James M. (2005) "How Probabilities Reflect Evidence," *Philosophical Perspectives* 19.1: 153-178.

Lempert, Robert J. (2002) "A New Decision Sciences For Complex Systems," *Proceedings of the National Academy of Sciences* 99.3: 7309-7313.

Lempert, Robert J. (2019) "Robust Decision Making (RDM)," in Marchau, Walker, Bloemen, and Popper, eds. *Decision Making under Deep Uncertainty: From Theory to Practice*, 23-29. Cham, CH: Springer Nature.

Lempert, Robert J. and Collins, Myles T. (2007) "Managing the Risk of Uncertain Threshold Responses: Comparison of Robust, Optimum, and Precautionary Approaches," *Risk Analysis* 27.4: 1009-1026.

Lempert, Robert J., Groves, David G., Popper, Steven W., and Bankes, Steve C. (2006) "A General Analytic Method for Generating Robust Strategies and Narrative Scenarios," *Management Science* 52.4: 514-528.

Lempert, Robert J., Popper, Steven W., and Bankes, Steven C. (2003) *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*. Santa Monica, CA: the RAND Corporation.

Levi, Isaac (1980) *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. Cambridge, MA: MIT Press.

Lewis, David K. (1980) "A Subjectivist's Guide to Objective Chance," in Jeffrey, ed. *Studies in Inductive Logic and Probability, Volume II*, 263-293. Berkeley, CA: University of California Press.

Luce, R. Duncan and Raiffa, Howard (1957) *Games and Decisions*: *Introduction & Critical Survey*. New York, NY: John Wiley & Sons.

Marchau, Walker, Bloemen, and Popper, eds. (2019) *Decision Making under Deep Uncertainty: From Theory to Practice*. Cham, CH: Springer Nature.

Ord, Toby (2020) *The Precipice: Existential Risk and the Future of Humanity*. London, UK: Bloomsbury.

Railton, Peter (1984) "Alienation, Consequentialism, and the Demands of Morality," *Philosophy & Public Affairs* 13.2: 134-171.

Robalino, David A. and Lempert, Robert J. (2000) "Carrots and Sticks For New Technology: Abating Greenhouse Gas Emissions in a Heterogeneous and Uncertain World," *Integrated Assessment* 1.1 (2000): 1-19.

Rosenhead, Jonathan, Elton, Martin, and Gupta, Shiv K. (1972) "Robustness and Optimality as Criteria for Strategic Decisions," *Operations Research Quarterly* 23.4: 413-431.

Sahlin, Nils-Eric (1993) "On Higher Order Beliefs," in Dubucs, ed. *Philosophy of Probability*, 13-34. Dordrecht, NL: Springer.

Savage, Leonard J. (1972) *The Foundations of Statistics, Second Revised Edition*. New York, NY: Dover Publications.

Schmidtz, David (2004) "Satisficing as Humanly Rational Strategy," in Byron, ed. *Satisficing and Maximizing*: *Moral Theorists on Practical Reason*, 30-58. Cambridge, UK: Cambridge University Press.

Schneller, G. O. and Sphicas, G. P. (1983) "Decision Making Under Uncertainty: Starr's Domain Criterion," *Theory and Decision* 15.4: 321-336.

Schwartz, Barry, Ben-Haim, Yakov, and Dasco, Cliff (2011) "What Makes a Good Decision? Robust Satisficing as a Normative Standard of Rational Decision Making," *Journal for the Theory of Social Behavior* 41.2: 209-227.

Schwartz, Peter (1996) *Art of the Long View: Planning for the Future in an Uncertain World*. New York, NY: Doubleday.

Selten, Reinhard (1998) "Aspiration Adaptation Theory," *Journal of Mathematical Psychology* 42.2-3: 191-214.

Simon, Herbert A. (1955) "A Behavioral Model of Rational Choice," *The Quarterly Journal of Economics* 69.1: 99–118.

Skyrms, Brian (1980) "Higher Order Degrees of Belief," in Mellor, ed. *Prospects for Pragmatism*, 109-137. Cambridge, UK: Cambridge University Press.

Slote, Michael (1989) *Beyond Optimizing: A Study of Rational Choice*. Cambridge, MA: Harvard University Press.

Slote, Michael (1984) "Satisficing Consequentialism," *Proceedings of the Aristotelian Society* 58.1: 139-163.

Sprenger, Jan (2012) "Environmental Risk Analysis: Robustness Is Essential for Precaution," *Philosophy of Science* 79.5: 881-892.

Starr, Martin K. (1966) "A Discussion of Some Normative Criteria for Decision-Making Under Uncertainty," *Industrial Management Review* 8.1: 71-78.

Tideman, T. N. (1987) "Independence of Clones as a Criterion for Voting Rules," *Social Choice and Welfare* 4: 185-206.

van Fraassen, Bas (1984) "Belief and the Will," *The Journal of Philosophy* 81.5: 235-256.

van Fraassen, Bas (1990) "Figures in a Probability Landscape," in Dunn and Gupta, eds. *Truth or Consequences*, 345-356. Dordrecht, NL: Springer.

Walker, Warren E., Lempert, Robert J., and Kwakkel, Jan H. (2013) "Deep Uncertainty," in Gass and Fu, eds. *Encyclopedia of Operations Research and Management*, *Third Edition*, 395-402. New York, NY: Springer Science + Business.

Weatherson, Brian (1998) "Decision Making With Imprecise Probabilities," Unpublished manuscript <http://brian.weatherson.org/vdt.pdf>