



Munich Personal RePEc Archive

Setting environmental policy when experts disagree

Athanassoglou, Stergios and Bosetti, Valentina

1 April 2012

Online at <https://mpra.ub.uni-muenchen.de/57494/>
MPRA Paper No. 57494, posted 26 Jul 2014 13:17 UTC

Setting environmental policy when experts disagree*

Stergios Athanassoglou and Valentina Bosetti[†]

February 2014; revised April 2014; June 2014

Abstract

How can a decision-maker assess the potential of environmental policies when a group of experts provides divergent estimates on their effectiveness? To address this question, we propose and analyze a variant of the well-studied α -maxmin model in decision theory. In our framework, and consistent to the paper's empirical focus on renewable-energy R&D investment, experts' subjective probability distributions are allowed to be action-dependent. In addition, the decision maker constrains the sets of priors to be considered via a parsimonious measure of their distance to a benchmark "average" distribution that grants equal weight to all experts. While our model is formally rooted in the decision-theoretic framework of Olszewski [37], it may also be viewed as a structured form of sensitivity analysis. We apply our framework to original data from a recent expert elicitation survey on solar energy. The analysis suggests that more aggressive investment in solar energy R&D is likely to yield significant dividends even, or rather especially, after taking expert ambiguity into account.

Keywords: expert elicitation; expert aggregation; decision theory; renewable energy R&D

JEL classifications: C60, D81, Q42, Q48

*We are grateful to Michela Catenacci and Giulia Fiorese for help with data collection and to Gauthier de Maere. We also thank seminar participants at FEEM and ETH Zurich, for useful comments. The research leading to these results has received funding from the European Research Council under the European Commission's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement #240895, project ICARUS. The views expressed herein are purely those of the authors and may not in any circumstances be regarded as stating an official position of the European Commission.

[†]Athanassoglou: Corresponding Author. Researcher, European Commission Joint Research Center, Econometrics and Applied Statistics Unit, Via Enrico Fermi, 2749, 21027 Ispra Varese, Italy, athanassoglou@gmail.com. Bosetti: Department of Economics, IEFE and IGIER, Bocconi University, Ettore Bocconi Department of Economics Room 5-D1-01, via Roentgen, 1, 20136 Milan, Italy, valentina.bosetti@unibocconi.it.

1 Introduction

Environmental policy makers must often come to grips with considerable uncertainty. The inherent complexity of many environmental problems, combined with a lack of reliable data, impose significant challenges for the analysis of different policy options. In such situations, guidance can be sought by directly consulting a group of knowledgeable experts.

The field of structured expert judgment (also known as “expert elicitation”) has been developed to add scientific rigor to the process of expert consultation. Expert elicitation derives probabilistic input for decision problems through the quantification of experts’ subjective uncertainties (Morgan and Henrion [31], Cooke [13], O’Hagan et al. [36]). Subjective probability distributions are derived via transparent protocols and treated as scientific data (Cooke [13]). The employed techniques involve recognizing and removing, as much as possible, known psychological biases in judgment (Tversky and Kahneman [42, 43]). Consistency checks are incorporated throughout, and questionnaires are designed to elicit clear, unequivocal answers. Recognizing the value-added that it brings to the decision making process, expert elicitation is used in a wide variety of settings and the interested reader may consult [31, 13, 36, 44] for a comprehensive account.

Expert elicitation studies are being increasingly pursued in the assessment of environmental policy. In the United States, they have become part of the government’s analytic toolkit (U.S. EPA [44]) in the development of environmental regulation. Their acceptance by policy makers no doubt reflects their growing popularity in academic circles. In the academic literature on climate change policy, numerous expert elicitations have been performed to provide estimates of critical climate-model parameters (e.g., climate sensitivity) as well as climate-related impacts (Nordhaus [34], Morgan et al. [32, 33], Vaughn and Spouge [46], Zickfeld et al. [49, 50]). Furthermore, and related to the empirical focus of the present paper, a number of recent studies have applied expert elicitation techniques to assess the potential of investment in renewable-energy R&D. Baker et al. [6, 7], Baker and Keisler [8], Curtright et al. [14], Bosetti et al. [10] and Fiorese et al. [16] collected and interpreted expert opinions to investigate the uncertain effects of R&D investments on the prospect of success of carbon capture and storage, hybrid electric vehicles, cellulosic biofuels, and solar PV technologies.

Despite its compelling features, expert elicitation often generates divergent opinions across experts, resulting in ambiguous recommendations. To address this inconclusiveness, researchers

typically aggregate over expert estimates in some fashion and consider the resulting weighted average in their subsequent assessment. Indeed, there is a rich, primarily management-science, literature that studies the many different ways such aggregations may be performed (Clemen and Winkler [11, 12]).

Yet, we can imagine settings in which the quest for a unique probability distribution summarizing expert opinion is either infeasible or undesirable. This is particularly true when the issue on which expert opinion is sought is unprecedented and/or exceedingly complex, implying an extremely wide range of estimates. In such instances, aggregating over experts and making decisions based on the resulting average probability distribution may lead to suboptimal policy. In environmental applications, this point has been forcefully made with regard to climate change (Knutti et al. [25], Millner et al. [30], Woodward and Bishop [47]) and ecosystem management (Ludwig et al. [27], Vardas and Xepapadeas [45]). In such cases, the decision maker may arguably be better off simultaneously exploring the consequences of many different weighting schemes in a systematic fashion.

A modeling framework for addressing such situations can be found in the literature on decision making under ambiguity.¹ In contrast to the traditional Bayesian setting, in which probabilities are assigned to events via a unique Bayesian prior, decision-theoretic models of ambiguity are designed to address situations in which a decision maker is unable to assign precise probabilistic structure to the problem she is trying to address. In such situations, the traditional expected utility paradigm (based on a unique prior summarizing all problem uncertainty) fails on both normative and descriptive grounds (Gilboa and Marinacci [21]). On normative grounds it clashes with appealing axioms of rational choice, while from a descriptive standpoint it is often in conflict with observed patterns of behavior.² Given the complexity of many environmental phenomena, it is unsurprising that models of ambiguity aversion are increasingly being employed to address policy under conflicting expert opinion (Woodward and Bishop [47], Millner et al. [30], Asano [3], Athanassoglou and Xepapadeas [4]).³

¹See Gilboa and Marinacci [21] for a comprehensive recent survey.

²Indeed, there seems to be experimental evidence suggesting that lab subjects actually reason in a way that is consistent to models of ambiguity aversion when confronted with inconclusive expert judgments (Baillon et al. [5]).

³An important additional concern that some authors have raised is that probabilistic estimates elicited from experts may not always represent probabilities in the actuarial sense (see, e.g., Millner et al. [29] and Kriegler et al. [26]). This concern seems to apply to the empirical application of this paper as well, at least in part (see Appendix

Our contribution. In this paper we further the literature on environmental decision-making under conflicting expert opinion. To fix ideas and prepare the ground for the paper’s empirical exercise, let us suppose that a decision maker elicits the judgment of a set of experts on the effect of R&D investment on the future cost of a promising green technology.⁴ Levels of R&D investment affect the decision maker’s problem in two ways: (a) they alter experts’ subjective probability distributions on the technology’s future cost and (b) they are arguments of a utility function that measures the investment’s cost-effectiveness as a function of R&D expenditure and the resulting technological improvement.

As experts will frequently provide divergent estimates on the potential of green R&D, the decision maker wishes to systematically explore the implications of many different ways of weighting their estimates. We propose a parsimonious parametric framework to facilitate this process. As an initial benchmark, our model posits equal-weight linear aggregation over experts’ divergent probability distributions. Subsequently, it considers enlargements of the set of possible linear aggregation schemes by parameterizing over their maximum distance, measured via the Euclidean norm, with respect to the benchmark equal-weight aggregation. This distance is referred to as *aggregation ambiguity* and can be naturally interpreted as setting a bound on the total weight that can be assigned to any group of experts. Guaranteeing a bound of this sort may be a normative or operational desideratum, in the sense that it provides a check on the influence of any single group of experts, and guards against outlier opinions. Next, our model computes the best-and worst-case expected outcomes of a given level of R&D investment, subject to the feasible set of distributions that is implied by assigned levels of aggregation ambiguity. Finally, the decision maker’s preferences are captured by a convex combination of the best and worst-case expected outcomes. The weights placed on the worst and best cases reflect her *ambiguity attitude*.

The above framework is a variation of the α -maxmin model that has been studied extensively in the economic-theory literature beginning with Arrow and Hurwicz [2]. Later contributions by Gilboa and Schmeidler [22] (whose seminal paper dealt with the pure maxmin model), Ghirardato

A1). The decision-theoretic literature has begun to address such imprecise probabilities (see, e.g., Gajdos et al. [17] and Gajdos and Vergnaud [18]) in the context of decision making under ambiguity. Extending our framework to account for expert imprecision would be a very interesting avenue of future research.

⁴Note, however, that our formal model generalizes to many alternative contexts of decision making under uncertainty.

et al. [20], and Eichberger et al. [15] provided axiomatic treatments of similar models. The model we adopt in this paper is not a strict application of this framework. This is because its decision variables enter the value function as arguments of both (a) a utility function measuring the technology’s payoff as well as (b) the set of priors that the decision-maker is taking into account when performing her best- and worst-case analysis. This latter element of action-dependent subjective beliefs is non-standard in the formal decision-theoretic literature going back to the work of Savage [40].⁵ Jaffray [23] had first introduced a similar notion with an α -maxmin model based on non-additive belief functions, while later Ghirardato [19] analyzed a model in which acts map from states to sets of consequences.

Recently, Olszewski [37] studied the α -maxmin model in a setting in which decision makers are called to choose between sets of lotteries, i.e., probability distributions over outcomes. This is the α -maxmin framework that most fits our context, since levels of R&D directly affect experts’ probability distributions –and thus can be seen as defining sets of lotteries over the technology’s future payoff. Olszewski showed that when choosing between sets of lotteries the α -maxmin rule uniquely satisfies a set of axioms that generalize their counterparts in settings of single probability distributions.

Our paper aims to operationalize the model axiomatized by Olszewski in an applied setting of expert elicitation. Indeed, our work is an outgrowth of the need to develop a tractable theoretical framework to accommodate expert opinions gathered by a recent elicitation [10]. Still, it is important to note that our framework is open to legitimate criticism. First, not everyone may agree with our consideration of aggregation ambiguity via the Euclidean norm. Second, the axiomatic underpinnings of the Olszewski α -maxmin model, though clear, will not, in general, be universally accepted as a means of characterizing the “right” decision maker preferences under ambiguity. Thus, even for researchers who espouse models of ambiguity aversion, our framework may not be entirely satisfactory.

On a more general level, a number of scholars are skeptical of using models of ambiguity aversion as a foundation for rational decision making. The application area of expert elicitation only serves to heighten these concerns, as specialized researchers have developed a variety of methods in order to arrive to a unique distribution representing expert beliefs (as discussed in the extensive survey

⁵Note, however, that endogenous probabilities are common in Decision Analysis and dynamic programming.

papers by Clemen and Winkler). We do not take a strong stand on this issue, not least because we ourselves are not philosophers of science, nor specialists in expert elicitation. Still, what we can constructively say is that for scholars who are unconvinced by models of ambiguity aversion, our framework may be interpreted as a kind of structured sensitivity analysis. The logic would run along the following lines:⁶ It is resource-intensive and perhaps politically sensitive to determine an optimal weighting of the experts. Thus, it is preferable to assume equal weights and then apply the machinery of the proposed model to test the robustness of the resulting recommendations. If we see that optimal R&D levels are highly sensitive to small changes in the model parameters, then we can conclude that assigning equal weights to experts may lead to non-robust policy and that more careful deliberation regarding expert weighting is called for.

The paper’s empirical section applies our model to original data from the ICARUS project (Bosetti et al. [10]), a recent expert elicitation on the potential of European Union R&D investment in renewable energy technologies (for more information see *www.icarus-project.org*). As an initial step, we used the collected data of the ICARUS survey to construct experts’ subjective probability distributions on the future cost of solar energy conditional on R&D investment. Subsequently, we employed an integrated assessment model (Bosetti et al. [9]) to calculate the benefits of R&D investment (via the future solar-electricity costs that R&D may bring about). The application of our theoretical model to these data suggested that expert ambiguity plays an important role in assessing the potential of solar technology. Our analysis allowed us to (cautiously) draw two policy implications: (1) that a 100% increase of EU investment in solar technology R&D is likely to yield significant rewards, even after taking expert ambiguity into account, (2) that a 50% increase in investment will likely always be less preferred to either a 100% increase, or maintaining the status quo.

Related work. There is a long and growing literature of applications of models of ambiguity aversion to environmental economics (see [45, 3, 41, 4, 30] among others). Most of these papers deal with different, dynamic models of ambiguity aversion, which do not have action-dependent subjective probabilities. Moreover, with the exception of Millner et al. [30], these papers do not engage directly with the issue of expert aggregation. Perhaps the closest parallel to our work can be found in Woodward and Bishop [47]. There, the authors applied a much simplified version of

⁶We thank an anonymous referee for suggesting this alternative interpretation of our work.

our α -maxmin framework to William Nordhaus’ DICE integrated assessment model [35]. There were two stylized “experts” in the analysis, reflecting different views on the model’s damage function parameters.⁷ The first expert held that damages will be relatively moderate (in line with Nordhaus’ “base-case scenario”), while the second was much more pessimistic (in line with Nordhaus’ “catastrophic” scenario). In this binary setting, the authors argued that a rational decision maker would seek to maximize the worst-case payoff, in the Gilboa and Schmeidler [22] sense, and adopt an emissions trajectory that is optimal under the catastrophic scenario. This finding was reinforced by the possibility of learning the true damage parameters at a later stage. While the analysis was quite simplified, the main policy implication was that it pays to follow a policy of aggressive abatement from the outset, even if the future is not quite as catastrophic as one might fear.

Paper outline. The structure of the paper is as follows. Section 2 discusses an example of our setting to motivate the introduction of the formal model, which follows in Section 3. Section 4 applies the theoretical model to original data from the ICARUS expert elicitation survey on solar technology and discusses possible policy implications. Section 5 provides brief concluding remarks and directions for future research. An Appendix collects non-essential supplementary information, as well as tables and figures not appearing in the main text.

2 Model description

Consider a set \mathcal{N} of experts indexed by $n = 1, 2, \dots, N$. R&D investment is denoted by a variable $r \in \mathcal{R}$ and the technology’s cost by $c \in \mathcal{C}$, where \mathcal{R} and \mathcal{C} are subsets of real numbers. To make our framework completely consistent with the model of Olszewki [37] we assume the sets \mathcal{R} and \mathcal{C} are finite. An expert n ’s probability distribution of the future cost of technology given investment r is captured by a random variable having a probability distribution function (pdf)

$$\pi_n(c|r). \tag{1}$$

How do we make sense of divergence in the pdfs $(\pi_1(\cdot|r), \pi_2(\cdot|r), \dots, \pi_N(\cdot|r))$ when studying optimal R&D investment? In the absence of data that could lend greater credibility to one expert

⁷In contrast to the more general setting of our paper, these uncertain parameters were exogenous to the choice of policy.

over another, one straightforward way would be to simply aggregate over all pdfs π_n as given by Eq. (1), so that we obtain an “aggregate” pdf $\bar{\pi}$, where

$$\bar{\pi}(\cdot|r) = \sum_{n=1}^N \frac{1}{N} \pi_n(\cdot|r). \quad (2)$$

This approach inherently assumes that each and every expert is equally likely to represent reality, and makes use of simple linear aggregation. However, this way of aggregating expert pdfs may not always result in robust recommendations.

We thus move beyond simple averaging. In our framework each expert n 's pdf $\pi_n(c|r)$ is weighted by the decision maker through a second-order probability p_n . The set of admissible second-order distributions \mathbf{p} depends on the amount of ambiguity the decision maker is willing to take into account when aggregating across experts, and in particular on how “far” she is prepared to stray from equal-weight aggregation. Specifically, we consider the set of second-order distributions $P(b)$ over a set of N experts, parametrized by $b \in [0, \frac{N-1}{N}]$ where

$$P(b) = \left\{ \mathbf{p} \in \mathfrak{R}^N : \mathbf{p} \geq \mathbf{0}, \sum_{n=1}^N p_n = 1, \left[\sum_{n=1}^N \left(p_n - \frac{1}{N} \right)^2 \right]^{\frac{1}{2}} \leq b \right\}. \quad (3)$$

Here, the set $P(b)$ captures the ambiguity of the decision-maker’s aggregation scheme in a parsimonious parametric manner; as such, we refer to parameter b as *aggregation ambiguity*.⁸ It can be naturally interpreted as setting a bound on the total weight that can be assigned to any group of experts (see following discussion on page 10). Guaranteeing a bound of this sort may be a normative or operational desideratum, in the sense that it provides a check on the influence of any single group of experts, and guards against outlier opinions. Letting \mathbf{e}_N denote a unit vector of dimension N , we see that distributions \mathbf{p} belonging to $P(b)$ satisfy $\|\mathbf{p} - \frac{\mathbf{e}_N}{N}\|_2 \leq b$, where $\|\cdot\|_2$ denotes the Euclidean norm. Setting $b = 0$ implies the unique consideration of equal second-order probabilities, while $b = \sqrt{\frac{N-1}{N}}$ consideration of the entire simplex $\Delta^{N-1} = \left\{ \mathbf{p} \in \mathfrak{R}^N : \mathbf{p} \geq \mathbf{0}, \sum_{n=1}^N p_n = 1 \right\}$.⁹

⁸Note that, despite the term “ambiguity”, the parameter b does not characterize a state or amount of knowledge. Knowledge, as represented by the set of experts’ pdfs, remains fixed.

⁹The latter statement holds in light of the fact that values of $b > \sqrt{\frac{N-1}{N}}$ cannot enlarge the feasible set. This is because the maximizers of $\left[\sum_{n=1}^N \left(p_n - \frac{1}{N} \right)^2 \right]^{\frac{1}{2}}$ over the set of probability vectors concentrate all probability mass on one expert, leading to an aggregation ambiguity of $\left[\left(1 - \frac{1}{N} \right)^2 + (N-1) \cdot \left(\frac{1}{N} \right)^2 \right]^{\frac{1}{2}} = \sqrt{\frac{N-1}{N}}$.

Fixing a level of R&D r and weighting the expert pdfs (1) under all linear aggregation schemes belonging in $P(b)$ induces the following set of probability distributions

$$\Pi(b, r) = \left\{ \sum_{n=1}^N p_n \pi_n(\cdot|r) : \mathbf{p} \in P(b) \right\}, \quad (4)$$

governing the future cost of the technology, conditional on R&D investment r . Thus, the set of distributions $\Pi(b, r)$ is identical to the set of mixtures of $\{\pi_1(\cdot|r), \pi_2(\cdot|r), \dots, \pi_N(\cdot|r)\}$, where the weights of the mixture distributions are constrained to lie in set $P(b)$. Consequently, holding r fixed, an increase in b implies an expansion of the set of priors a decision maker is willing to consider.¹⁰

Now, define the real-valued function

$$u(c, r) : \mathcal{C} \times \mathcal{R} \mapsto \mathfrak{R},$$

as representing the utility of R&D investment r , under cost realization c . Given choices for r and b , we can calculate the best- and worst-case expected utilities associated with r , given the set of induced prior distributions $\Pi(b, r)$. This provides a measure of the spread between the worst- and best-cases, given a “willingness” to stray from the benchmark equal-weight distribution (2) that is constrained by an aggregation ambiguity of b . More formally, we consider the functions

$$V_{min}(r|b) = \min_{\pi \in \Pi(b, r)} \int_{\mathcal{C}} u(c, r) d\pi(c) \quad (5)$$

$$V_{max}(r|b) = \max_{\pi \in \Pi(b, r)} \int_{\mathcal{C}} u(c, r) d\pi(c). \quad (6)$$

Plotting functions (5) and (6) over $b \in [0, \sqrt{(N-1)/N}]$ gives decision makers a comprehensive picture of the expected payoff of R&D investment r .

¹⁰It is important to note that the parameterization (4) may, in some cases, obscure important information. For instance, suppose one wishes to make a decision on the basis of a water container’s temperature. There are two experts, one of which provides a pdf suggesting that the water is freezing and the other that it is boiling hot. Suppose, further, that we know that one of the two experts *must* be correct. Then it is clear that if we consider the set of pdfs (4) for some small value of b , we will be taking into account a whole set of pdfs implying that the water is tepid and none that it is either freezing or boiling. Clearly, in such cases, our framework will likely lead to suboptimal decisions. To avoid this from happening, we should either (a) set $b \approx \sqrt{\frac{1}{2}}$ that allows for the unmixed consideration of both expert pdfs, or (b) consider a different parametric framework altogether. We thank an anonymous referee for bringing this issue to our attention.

Thus, given a level of aggregation ambiguity b , the functions (5)-(6) focus on the worst- and best-case expected utility associated with an investment r . As such they capture extreme attitudes towards uncertainty in expert aggregation. To express more nuanced decision-maker preferences we consider the following value function

$$V(r|b, \alpha) = \alpha \cdot V_{min}(r|b) + (1 - \alpha) \cdot V_{max}(r|b) \quad \alpha \in [0, 1], \quad (7)$$

representing a convex combination of the worst- and best-case expected utility. The parameter α above captures the decision maker's *ambiguity attitude*. It measures her degree of pessimism given aggregation ambiguity b : the greater (smaller) α is, the more (less) weight is placed on the worst-case scenario. In Section 3, we explain more formally the sense in which parameter α captures the decision maker's attitude towards ambiguity.

As mentioned in the introduction, Eq. (7) is a variation of the well-studied α -maxmin model in decision theory. Given values for b and α , Eq. (7) operates as an objective function when searching for optimal R&D investment r .

What do different choices of b imply? We offer some brief remarks on the interpretation of different levels of aggregation ambiguity b . Consider the benchmark equal-weight aggregation $\frac{1}{N} \mathbf{e}_N$. Now take any subset of experts $\hat{\mathcal{N}}$ of cardinality \hat{N} and begin increasing the collective second-order probability attached to their pdfs. The convex structure of the feasible set $P(b)$ enables us to provide a tight upper bound on the maximum total second-order probability that can be placed on this set of experts, as a function of b and \hat{N} (we denote $|\hat{\mathcal{N}}| = \hat{N}$):

$$\max_{\mathbf{p} \in P(b)} \sum_{n \in \hat{\mathcal{N}}} p_n = \min \left\{ \frac{\hat{N}}{N} + b\hat{N} \sqrt{\frac{N - \hat{N}}{\hat{N}N}}, 1 \right\}. \quad (8)$$

Extending this logic to any subset of experts, we have the following holding:

$$\mathbf{p} \in P(b) \Leftrightarrow \left\{ \mathbf{p} \geq \mathbf{0}, \sum_{n=1}^N p_n = 1, \sum_{n \in \hat{\mathcal{N}}} p_n \leq \min \left\{ \frac{\hat{N}}{N} + b\hat{N} \sqrt{\frac{N - \hat{N}}{\hat{N}N}}, 1 \right\}, \text{ for all } \hat{\mathcal{N}} \subseteq \mathcal{N} \right\}. \quad (9)$$

Figure 1 provides a graphical visualization of the above algebraic arguments for the case of $N = 3$ experts. Consider two levels of aggregation ambiguity b_1 and b_2 . Assume that b_1 is such that $\frac{2}{3} + b_1 * 2\sqrt{\frac{3-2}{2*3}} < 1 \Leftrightarrow b_1 < \sqrt{\frac{1}{6}}$. In this case, by Eq. (8) we know that no pair of experts can




Figure 1: The set $P(b)$ for $N = 3$ and two levels of aggregation ambiguity $b_1 < \sqrt{\frac{1}{6}} < b_2 < \sqrt{\frac{2}{3}}$. $P(b_1)$ is the inner circle of radius b_1 , while $P(b_2)$ is the oval-shaped triangular figure subsuming $P(b_1)$.

be assigned weight of 1, implying that all experts must be assigned positive weight. Hence, the set $P(b_1)$ is just a circle of radius b_1 that is embedded in the simplex Δ^2 and centered at the equal-weight benchmark $(1/3, 1/3, 1/3)$. This circle does not touch the simplex's boundary. Conversely, when $b = b_2$ such that $b_2 > \sqrt{\frac{1}{6}}$ and $\frac{1}{3} + b_2 \sqrt{\frac{3-1}{1*3}} < 1 \Leftrightarrow b_2 < \sqrt{\frac{2}{3}}$, the decision maker (a) allows for a weight of 1 to be placed on sets of pairs of experts (and thus 0 for sets of single experts), while (b) setting an upper bound of $\frac{1}{3} + b_2 \sqrt{\frac{2}{3}} < 1$ on the weight than a single expert can receive. In this case, the set $P(b_2)$ has a triangular oval-like shape, with parts of it intersecting the boundary of Δ^2 . Moreover, the closer b_2 is to $\sqrt{\frac{2}{3}}$, the closer the curved segments of $P(b_2)$ are to becoming two intersecting straight lines. When $b = \sqrt{\frac{2}{3}}$, this transformation is complete and we have $P(b) = \Delta^2$.

On a final note we recognize that, when applying the framework of Eq. (7), there may not be one single objectively “correct” value for b and α . Thus, we believe it is preferable to examine Eq. (7) over a reasonable range of b (guided by the insights of Eq. (9)) and α , that is deemed appropriate for the context at hand. Such an investigation will allow the decision maker to explore the implications of different (b, α) -combinations in a systematic fashion, and ultimately arrive at a

choice of r that she considers justified.¹¹

3 Decision-theoretic foundations

In this section we provide an exposition of the axiomatic underpinnings of the decision-theoretic model we adopt in Eq. (7). In doing so, we use the analytic framework of Olszewski [37], which is the most appropriate for our context. Olszewski axiomatized a version of the α -maxmin model in a setting in which decision makers are called to choose between sets of lotteries, i.e., probability distributions over outcomes. This fits our R&D setting since, in our model, levels of R&D directly affect experts' probability distributions, as per Eq. (1), and thus can be seen as defining a set of lotteries over the technology's future payoff.

Readers familiar with models of ambiguity aversion may wonder why we do not draw on Ghirardato et al. [20] to discuss the axiomatic foundations of our α -maxmin model. More broadly, they might also ask why we do not adopt different models of ambiguity aversion such as the smooth ambiguity model of Klibanoff et al. [24]. The short answer to these questions is that the endogeneity of expert probabilities of Eq. (1) distinguishes our framework from the decision-theoretic setting of Savage [40], on which the previous contributions are grounded. In particular, the preferences analyzed in [20, 24] operate in settings where ambiguous probabilistic information is independent of a decision maker's actions, and are thus not directly applicable to our context.¹²

To further highlight the above point, it is instructive to compare our setting to that of Millner et al. [30]. The starting point of Millner et al.'s analysis is that different climate models (representing different "expert" groups of climate scientists) provide widely divergent estimates of the probability distribution of climate sensitivity, a key physical parameter linking greenhouse gases and temperature increase (see Meinshausen et al. [28]). The scientific community has not reached an agreement on how best to aggregate these divergent estimates and arrive to a single pdf of climate sensitivity. Acknowledging this lack of scientific consensus on the "correct" pdf of climate sensitivity, Millner et al. [30] adopted the smooth ambiguity model of Klibanoff et al. [24] to study optimal abatement policy. However, a key aspect of their modeling framework is that nothing a

¹¹The empirical example of Section 4 (especially Figure 4) suggests the logic of such an exercise.

¹²On a related note, the theoretical machinery of the Savage framework (involving state spaces, consequence spaces, and acts) seems overwrought for our context (to this effect, see also the discussion in the introduction of Ahn [1]).

decision maker does can alter ambiguous scientific information; no decision on abatement can affect scientists' probabilistic estimates of climate sensitivity. Climate sensitivity is a physical quantity that is simply unaffected by human action. This is clearly in contrast to the setting of the present paper in which decisions on R&D investment do affect experts' probabilistic estimates on the future cost of the green technology via Eq. (1).

In what follows, we adopt the notational conventions of Olszewski [37]. Suppose $X = \{x_1, x_2, \dots, x_n\}$ is a finite set of outcomes, and let $\Delta = \Delta(X)$ be the set of all lotteries (i.e., probability distributions) on X . Moreover, let \mathcal{A} denote the set of all closed and non-empty subsets of Δ . Given a utility function $u : X \mapsto \mathfrak{R}$ and a lottery $l = (q_1, q_2, \dots, q_n)$, where q_i denotes the probability of outcome x_i , let

$$U(l) = \sum_{i=1}^n q_i u(x_i)$$

denote the expected utility of the lottery l . Consider now a complete and transitive preference relation \preceq over \mathcal{A} . For the rest of this chapter the following assumption is assumed to hold.

Assumption 1 *The preferences \preceq restricted to single lotteries admit an expected utility representation. That is, there exists a utility function $u : X \mapsto \mathfrak{R}$ such that for any two lotteries $l_1, l_2 \in \Delta$, $l_1 \preceq l_2 \Leftrightarrow U(l_1) \leq U(l_2)$.*

As usual, the utility function u captures the decision maker's attitude towards risk in an expected utility framework. Now, consider the α -maxmin rule

$$H_{u,\alpha}(A) = \alpha \left(\min_{l \in A} U(l) \right) + (1 - \alpha) \left(\max_{l \in A} U(l) \right), \quad (10)$$

where $\alpha \in [0, 1]$, and define the corresponding preferences $\preceq^{H_{u,\alpha}}$ on \mathcal{A} , whereby

$$A_1 \preceq^{H_{u,\alpha}} A_2 \Leftrightarrow H_{u,\alpha}(A_1) \leq H_{u,\alpha}(A_2).$$

Evidently, the value function we consider in Eq. (7) reduces to a special case of the preferences defined in Eq. (10). We briefly explain how. In the setting of Section 2, the relevant outcome space equals the cross product of the cost and R&D domains, i.e. $\mathcal{C} \times \mathcal{R}$. For an expert n , a choice of r can be viewed as defining a pdf π_n on $\mathcal{C} \times \mathcal{R}$, such that $\pi_n(\cdot, r) \equiv \pi_n(\cdot|r)$ (where the latter is given by Eq. 1) and $\pi_n(\cdot, \tilde{r}) \equiv 0$ for all $\tilde{r} \neq r$.¹³ Consequently, given a (fixed) value of aggregation

¹³Note that the presence of r as an argument of utility function u compels us to consider this broader outcome space (instead of just \mathcal{C}) and somewhat awkward reframing of pdfs π_n , even though there is no direct randomness associated with r (once the decision maker chooses it, it is fixed).

ambiguity b , choosing a value of r is equivalent to choosing a set of lotteries $\Pi(b, r) \in \mathcal{A}$, over which the minimum and maximum expected utilities will be computed.

Olszewski’s article performed a rigorous and comprehensive theoretical analysis of preferences $\preceq^{H_{u,\alpha}}$. In what follows we selectively present what is necessary for our purposes. We are primarily interested in shedding light on two main issues: (i) the axiomatic foundations of rule $H_{u,\alpha}(A)$, and (ii) the interpretation of comparative attitudes towards ambiguity.

Axiom 1 [Set S-Solvability]. *For all triplets of sets A, A_1, A_2 such that one of the sets A_1, A_2 is a singleton, if $A_1 \preceq A \preceq A_2$, then there exists a $p \in [0, 1]$ such that $A \sim pA_1 + (1 - p)A_2$.*

As we can see, Set S-Solvability operates as an analog of the intermediate value theorem of calculus on preferences over sets of lotteries. As such, it imposes a sort of continuity property on preferences so that no “jumps” are allowed.

A stronger, and more mathematically involved, version of preference continuity is given by the following axiom.

Axiom 2 [Set Continuity]. *For every set of lotteries A_1 , sets $\{A_2 : A_2 \prec A_1\}$ and $\{A_2 : A_1 \prec A_2\}$ are open in the space of all closed subsets of Δ equipped with the Hausdorff distance.¹⁴*

Set Continuity implies Set S-Solvability, and generalizes the traditional notion of continuity to sets of lotteries. Among its consequences is the fact that a decision-maker will be indifferent between a set of lotteries A and $A \cup l$, where $l \notin A$ is a “middle-of-the-line” lottery that we know will never be among the argmins or the argmaxs in Eq. (10). This is a strong, and not uncontroversial, axiom. See Chapter 4 in Olszewski [37] for more information on how Set Continuity may, in some cases, be violated by decision makers.

Axiom 3 [Set S-Independence]. *If $A_1 \preceq A_2$ and l is a single lottery, then $pA_1 + (1 - p)l \preceq pA_2 + (1 - p)l$ for all $p \in [0, 1]$.*

¹⁴Skipping some of the formal details, the Hausdorff distance of two lotteries A and B is given by

$$d_H(A, B) = \max \left\{ \max_{a \in A} \min_{b \in B} d(a, b), \max_{b \in B} \min_{a \in A} d(a, b) \right\}.$$

Set S-Independence generalizes the traditional notion of weak independence for single lotteries. Essentially, what it says is that the decision maker's preferences over two sets A_1 and A_2 cannot be reversed by randomizing over these sets and a single lottery l .

Axiom 4 [Disjoint Set Betweenness (DSB)]

- (a) For any disjoint sets A_1 and A_2 , if $A_1 \preceq A_2$, then $A_1 \preceq A_1 \cup A_2 \preceq A_2$.
- (b) If $l_1 \prec l_2$ for all $l_1 \in A_1$ and $l_2 \in A_2$, then $A_1 \prec A_1 \cup A_2 \prec A_2$.

Part (a) of DSB captures what seems like an intuitive property: if two sets of lotteries A_1 and A_2 are disjoint and A_2 is preferred to A_1 , then their union, in terms of the decision maker's preferences, will lie somewhere inbetween A_1 and A_2 . That is, one cannot gain by incorporating the “bad” lotteries of A_1 to A_2 , in the same way as one cannot lose by incorporating the “good” lotteries of A_2 to A_1 . Part (b) restates (a) for the case of disjoint sets of lotteries in which the lotteries of one set uniformly dominate the lotteries of the other.

Theorem 1 (Olszewski [37]) *A preference \preceq defined on \mathcal{A} satisfies Set Continuity, Set S-Independence, and DSB if and only if there exists a unique $\alpha \in (0, 1)$ such that*

$$\preceq = \preceq^{H_{u,\alpha}} .$$

The parameter α is unaltered under positive affine transformations of u . As a corollary, preferences $\preceq^{H_{u,\alpha}}$ satisfy Set S-Solvability, Set S-Independence, and DSB for any $\alpha \in (0, 1)$.

Having presented the characterization result of Theorem 1, we now discuss comparative attitudes towards ambiguity. Given preferences \preceq^1 and \preceq^2 , let B be a set of lotteries such that \preceq^1 and \preceq^2 coincide on single lotteries from the set B . Then B is called *congruent* with respect to preferences \preceq^1 and \preceq^2 . More formally, B is congruent with respect to \preceq^1 and \preceq^2 if for all $l_1, l_2 \in B$, $l_1 \preceq^1 l_2 \Leftrightarrow l_1 \preceq^2 l_2$.

Definition 1 *Consider two preferences \preceq^1 and \preceq^2 . If, for any congruent set B ,*

$$B \preceq^2 A \Rightarrow B \preceq^1 A \text{ and } B \prec^2 A \Rightarrow B \prec^1 A,$$

for all $A \subset B$, then a decision maker with preferences \preceq^1 is more ambiguity averse than one with \preceq^2 .

The intuition behind the above definition is the following. Preferences \preceq^1 are more ambiguity averse than \preceq^2 if, given any congruent set B , whenever \preceq^2 prefers a set of lotteries A to its bigger superset $B \supset A$, then so does \preceq^1 (while the other side of the relation may not always be true). This definition of relative ambiguity aversion is reminiscent of similar concepts in the literature of risk aversion (see, e.g., Yaari [48]). It leads to the following corollary, which identifies α as a measure of comparative ambiguity aversion.

Corollary 1 *Suppose that \preceq^1 and \preceq^2 defined on \mathcal{A} satisfy Set Continuity, Set S-Independence, and DSB. Then preferences \preceq^1 are more ambiguity averse than \preceq^2 if and only if $\alpha_1 \geq \alpha_2$.*

In closing this Section, we wish to note that Olszewski proved a variant of Theorem 1 and Corollary 1 for the restricted domain of convex polyhedral lotteries. This alternative characterization uses a different set of axioms (e.g., replacing Set Continuity), which may be useful for future analyses focusing on convex-polyhedral sets of lotteries.

4 Empirical application to solar-technology R&D

We base the empirical application of our model on original data collected by the ICARUS survey, an expert elicitation on the potential of solar technologies. During the course of 2010-2011, the ICARUS survey collected expert judgments on future costs and technological barriers of different Photovoltaic (PV) and Concentrated Solar Power (CSP) technologies.¹⁵ Sixteen leading European experts from academia, the private sector, and international institutions took part in the survey. The elicitation collected probabilistic information on (1) the year-2030 expected cost of the technologies; (2) the role of public European Union R&D investments in affecting those costs; and (3) the potential for the deployment of these technologies (both in OECD and non-OECD countries). We refer readers interested in the general findings of the survey to Bosetti et al. [10] and we focus here on the data on future costs as they form the basis of our analysis.

Current 5-year EU R&D investment in solar technology is estimated at 165 million US dollars. The ICARUS study elicited the probabilistic estimates of the 16 experts on the 2030 solar electricity

¹⁵The survey is part of a 3-year ERC-funded project on innovation in carbon-free technologies (ICARUS - Innovation for Climate chAnge mitigation: a study of energy R&D, its Uncertain effectiveness and Spillovers www.icarus-project.org).

cost (2005 c\$/kWh) under three future *Scenarios*: (1) keeping current levels of R&D constant until 2030, (2) increasing them by 50%, and (3) increasing them by 100%. Coherent responses were obtained from 14 out of the 16 experts so the analysis that follows focuses solely on them.¹⁶ We used linear interpolation of the survey’s collected data (generally 3-6 points of each expert’s cumulative distribution function(cdf) conditional on R&D investment) to compute a pdf for each expert $n \in \{1, 2, \dots, 14\}$, given the three relevant levels of R&D investment denoted by $r \in \{r_1, r_2, r_3\}$ (here r_i refers to Scenario i).¹⁷ These pdfs represent experts’ subjective probability distributions of the cost of technology as denoted in Eq. (1). Figure 3 plots the corresponding cdfs as well as the cdf that the aggregate pdf (2) leads to, under all three Scenarios.

[FIGURE 2 here]

As one can see in Figure 2 there is considerable disagreement between experts over the potential of solar technology. This disagreement is particularly acute under Scenario 1, and diminishes as R&D levels increase. Nonetheless, the breakthrough nature of innovation and the need to cross certain firm cost thresholds, means that ambiguity in expert estimates remains an important concern, even under Scenario 3. This will become apparent in the analysis to follow.

We measure the utility of an investment via its net payoff. Denoting the benefit associated to a technology cost c by the function $B(c)$ and the opportunity cost of an investment r by $O(r)$, this is given by the following utility function:

$$u(c, r) = B(c) - O(r). \tag{11}$$

The next section describes how we provide numerical values for $B(\cdot)$ and $O(\cdot)$.

Quantifying benefits and opportunity costs of solar technology R&D. Expected benefits of solar technology R&D investments are quantified via a general equilibrium intertemporal model

¹⁶By “coherent” we mean non-contradictory responses which do not violate the laws of probability. In particular, 2 out of the 16 elicited experts provided probabilistic estimates implying non-monotonic cumulative distribution functions, which is clearly in conflict with probability theory and thus unusable for our purposes. As mentioned in the introduction, such incoherent expert estimates have been discussed in the literature (Millner et al. [29], Kriegler et al. [26]), and addressing them in the context of expert elicitations is deserving of further study.

¹⁷Please refer to the Appendix for more information on how expert pdfs were constructed from the ICARUS survey data.

(Bosetti et al. [9]) that can account for a range of macro-economic feedbacks and interactions. These include the effects of energy and climate change policies, the competition for innovation resources with other power technologies, the effect of growth, as well as a number of other factors.¹⁸ To capture the long-term nature of such investments, the integrated assessment model is run over the time horizon 2005-2100 in 5-year time periods for the whole range of exogenously-imposed possible 2030-costs of solar power that we are considering. Subsequently, simulation results are compared to the benchmark case in which the cost of solar power is so high that the technology is not competitive with alternative production modes. For each possible 2030 solar-power cost, the benefit to the European Union is quantified by the discounted EU-consumption improvement over the entire time-horizon 2005-2100 with respect to the case where solar technology is not competitive. Table 1 summarizes the results.

[TABLE 1 here]

Three important assumptions are at the basis of the numbers reported in Table 2. First, as the survey concentrated on public EU R&D investment and the effects of increasing it, we disregard spillovers and technological transfers to the rest of the world and consider only the consumption improvement for Europe. Second, we evaluate the benefit of alternative 2030 costs of solar power assuming that no carbon policy is in place and that no special constraints on other technologies are imposed (e.g., a partial ban on nuclear technology). Third, we discount cash flows using a 3% discount rate. Although our choice is well in the range of discount rates adopted for large scale public projects, it is important to note that the cost threshold for positive payoffs is robust for a wide range of more myopic discount rate values. Our assumptions all err on the side of being conservative about the potential payoffs of solar-technology R&D.

We now explain how we calculate the costs of solar R&D investment. Given an R&D investment r , we assume that actual R&D spending is fixed at r during the period 2005-2030, in line with the survey questions. After 2030 we assume that spending drops to half its initial value, i.e. $r/2$, and remains at that level until 2100. This drop occurs because we assume that post-2030 funds represent the government's commitment to maintain the technological gains achieved by 2030. We

¹⁸The analysis is carried out using the World Induced Technical Change Hybrid (WITCH) model (Bosetti et al. [9]), an energy-economy-climate model that has been used extensively for economic analysis of climate change policies. See www.witchmodel.org for a list of applications and papers.

now derive the discounted opportunity cost of such expenditure streams of solar-technology R&D spending. In doing so we follow Popp [38] and assume that, at every time period, this opportunity cost is equal to 4 times the original investment. Thus, in our model the opportunity cost of a level of R&D investment r is given by the net present value of the stream $O(t)$ where $O(t) = 4 \cdot r$ for $t = 1, 2, \dots, 6$ and $2 \cdot r$ for $t = 7, 8, \dots, 20$ (once again we use a 3% discount rate). Table 2 summarizes these results for the three R&D Scenarios that the ICARUS survey focused on.

[TABLE 2 here]

Application of the decision-theoretic framework. We now extend our analysis to explicitly account for aggregation ambiguity and adopt the decision-theoretic model introduced in Section 3.¹⁹ Our objective is to compare the three R&D Scenarios, and we do not consider optimizing over a wider R&D domain \mathcal{R} . We make this choice primarily because we wish to keep the applied section brief and pursue more in-depth empirical analysis in future work.²⁰

Figure 3 plots $V_{max}(r|b)$ and $V_{min}(r|b)$ over $b \in \left[0, \sqrt{\frac{13}{14}}\right] \approx [0, .96]$ for the three Scenarios.

[FIGURE 3 here]

Focusing first on Scenario 1, we note that equal-weight aggregation of expert opinion (corresponding to $b = 0$) yields an expected payoff of approximately $\$1.36 \times 10^9$. We observe that the worst-case expected payoff drops to about $\$ - 3.4 \times 10^9$ at $b \approx .25$ at which point it largely stops being sensitive to changes in b , slowly asymptoting to its minimum value of $\$ - 3.67 \times 10^9$; in contrast, the best-case one increases steadily to a maximum value of $\$22.7 \times 10^9$ at the maximum level of $b = .96$. Under Scenario 2, the expected payoff for $b = 0$ is equal to $\$7.8 \times 10^9$. Subsequently, we see that the worst-case payoff drops to 0 at $b = 0.15$, at which point it keeps decreasing at a smaller rate until it practically reaches its minimum value of $\$ - 5.5 \times 10^9$ at $b \approx .55$. Conversely, the best-case payoff rises steadily to about $\$32 \times 10^9$ for $b \approx .55$ at which point it continues to

¹⁹All simulations are performed in Mathematica. The optimizations are facilitated by the fact that right-hand-side of Eqs. (5) and (6) can be rewritten as $V_{min}(r|b) = \min_{\mathbf{p} \in \mathcal{P}(b)} \sum_{n=1}^N p_n u_n(r)$ and $V_{max}(r|b) = \max_{\mathbf{p} \in \mathcal{P}(b)} \sum_{n=1}^N p_n u_n(r)$, where $u_n(r) \equiv \int_{\mathcal{C}} u(c, r) d\pi_n(c|r)$. These equalities are valid by the linearity of the expectation operator under mixtures of distributions.

²⁰Indeed, constructing plausible approximations of experts' $u_n(r)$ functions over an interesting range of r will likely require further engagement with the experts.

rise at a much smaller rate until it reaches a maximum value of $\$33.3 \times 10^9$ at $b = .96$. Thus for Scenario 2, aggregation uncertainty becomes largely unimportant once b reaches the threshold of 0.55. Under Scenario 3 the unambiguous payoff is around $\$20 \times 10^9$, significantly higher than both other Scenarios. The worst-case expected payoff drops relatively smoothly to a minimum value of $\$ -7.35 \times 10^9$ for $b = .96$, while the best-case one rises at a comparatively higher rate to $\$70.9 \times 10^9$.

It is clear that aggregation ambiguity is important under Scenario 3, for both the worst- and best-case expected payoffs, resulting in a wider spread than Scenarios 1 and 2. This fact is interesting in light of Figure 2, which shows that experts' pdfs are much more dispersed under Scenarios 1 and 2 than they are under 3. The reason behind this seemingly unexpected result is straightforward. As Table 2 suggests, expected payoffs of R&D investment are very sensitive at low cost values, i.e., less than 8c\$/kWh. The more aggressive investment of Scenario 3 has a greater effect on these lower cost values, and therefore its best- and worst-case payoffs are in turn more sensitive to changes in b .

Figure 4 goes a step further and compares the three R&D Scenarios for all possible combinations of b and α . Following the color scheme of Figure 3, a region's color corresponds to the Scenario that performs the best within it, while the bold numbers within regions denote the ranking of the three Scenarios within this range of (b, α) (e.g., an expression "321" means Scenario 2 dominates 1, and Scenario 3 dominates both 2 and 1).

[FIGURE 4 here]

Figure 4 makes clear that Scenario 3 dominates 1 and 2 for an extremely wide range of combinations of b and α . Conversely, Scenario 1 is the best option for a combination of very high b and α . Somewhat surprisingly, we see that Scenario 2 is dominated by either 1 or 3 for all possible combinations of b and α and thus will never be chosen by a decision maker whose preferences are captured by Eq. (7). Thus, on the basis of the presented data, it is clear that policy makers should opt for the most aggressive R&D investment, unless they are both (a) open to ignoring a very large set of surveyed experts (b) extremely concerned about the possibility of worst-case failure. Moreover, assuming all three options are readily implementable, they can safely disregard the middle-range R&D investment implied by Scenario 2.

We end this section by noting that the result of Figure 4 is quite robust to the application of the model to different subsets of experts. This is due to, in large part, the fact that a significant

majority of experts agree that doubling current R&D has a good chance of pushing the cost of solar below 7.5, which, according to WITCH and our assumption regarding opportunity costs, is the threshold point for large gains to be realized. That being said, it also should be mentioned that the views of Experts 4, 6, and 9, are at complete odds with this picture, so that applying the model on their pdfs alone results in completely different ranking of the three scenarios in (b, α) -space. In particular, Scenario 2 predominantly dominates 1 and 3, and where that is not the case Scenario 1 is preferred (Scenario 3 is dominated everywhere by either 1 or 2). This suggests that the pattern of Figure 4 is not solely driven by assumptions embedded into the WITCH model and the calculation of opportunity costs.

5 Conclusions

Expert elicitation studies can play an important role in informing environmental policy. If designed well, they may be able to capture in a transparent and objective way subjective probabilities that can subsequently be used as scientific data in the decision maker’s deliberations.

When there is significant variation in expert estimates, condensing all of the problem’s uncertainty into a single probability distribution reflecting equal weights across experts may yield policy recommendations that are not robust, especially when the issue on which expert opinion is sought is unprecedented or highly complex. To address such concerns, we proposed and analyzed a variant of the well-studied α -maxmin model due to Olszewski [37]. In line with the paper’s focus on R&D investment, decision variables in our model affect experts’ subjective probability distributions of the future cost-effectiveness of an investment. We applied our framework to original data from a recent expert elicitation survey on solar technology. The analysis suggested that more aggressive investment in solar technology R&D is likely to yield substantial benefits even after ambiguity over expert opinion has been taken into account.

It should be mentioned that our work, while formally rooted in the literature of decision making under ambiguity, may also be interpreted as a structured form of sensitivity analysis in a purely Bayesian setting. Thus, it may be of interest also to scholars who are skeptical of models of ambiguity aversion as a foundation of rational decision making.

Our work suggests several fruitful avenues for future research. A particularly challenging one would be to extend the α -maxmin model to take into account higher moments or quantiles (see,

e.g., the quantile-based decision-theoretic framework of Rostek [39]) of the considered mixture distributions of Eq. (4). This change would introduce nonlinearities that significantly complicate the analysis, both theoretically and computationally. Alternatively, we could consider making expert weights quantile-specific, to reflect varying degrees of decision-maker confidence across the domain of expert distributions. For instance, a decision maker may trust some experts more on the probability of tail events, while others more on the probability of typical events, depending on their experience.²¹ On the applied front, extensions of this work could delve deeper into the ICARUS survey data to obtain more data points and ultimately solve for optimal R&D investment over a broader domain.

Appendix

A1: Constructing expert pdfs for the three R&D Scenarios from ICARUS survey data

In the ICARUS survey, experts were asked to provide values for the 10th, 50th, and 90th percentile of their distributions for the 2030 cost of solar technology conditional on all three Scenarios. In addition, they were asked to provide values for the probability of this cost being less than or equal to the following three values: 11.3, 5.5, and 3c\$/kWh. These “threshold” cost levels correspond to projections of the costs of electricity from fossil fuels or nuclear in 2030. The first (11.27 c\$/kWh) corresponds to the 2030 projected cost of electricity from traditional coal power plants in the presence of a specific policy to control CO₂ emissions (thus effectively doubling electricity costs from fossil sources). The second threshold cost (5.5 c\$/kWh) is the projected cost of electricity from traditional fossil fuels in 2030, without considering any carbon tax. Finally, the third (3 c\$/kWh) reflects a situation in which solar power becomes competitive with the levelized cost of electricity from nuclear power.

Asking experts the follow up question on the likelihood of reaching threshold cost targets allowed the survey authors to guard against the cognitive pitfalls associated with direct elicitation of subjective probabilities, to increase the amount of elicited information, and to deepen the discussion with the expert, hence improving their perception of his/her beliefs. In cases where the two sets of

²¹We thank an anonymous referee for this suggestion.

answers (percentile values and threshold probabilities) were inconsistent, we contacted the expert in order to obtain coherent estimates. Moreover, we asked all experts to give values for the upper and lower limits of their distribution’s support in order to pinpoint the intervals over which their implied probability distributions range.

Such corrected estimates were obtained from 14 out of the original 16 experts, and therefore the analysis of Section 4 focuses solely on them. Among the respondents, not all provided values on the left and right endpoints of their distributions’ support. As a result, we deduced between 6 and 8 points of 14 experts’ cumulative distribution functions (cdf) of the 2030 cost of solar electricity, given the aforementioned three R&D investment Scenarios. From these points a probability distribution function (pdf) was constructed using linear interpolation in the following way. First of all, and in accordance with the experts’ answers, we considered cost levels c lying in $[2c\$/kWh, 30c\$/kWh]$ and discretized this interval on a scale of 0.5 (30c\$/kWh represents an estimate of the technology’s current cost). Now, suppose an expert reported the values of his/her cdf F_n at two successive points c_1 and c_2 where $c_2 > c_1$ and gave no further information on cost levels between c_1 and c_2 . Assuming right-continuity of F_n we took the probability mass $F_n(c_2) - F_n(c_1)$ to be distributed uniformly among the cost levels $\{c_1 + .5, c_1 + 1, \dots, c_2\}$. For experts who did not provide values for the lower limit of their distribution’s support we assumed that whatever probability mass remained to be allocated (always less than .1) was distributed uniformly between the smallest argument of the cdf and two cost levels below it. For example, if an expert only indicated that c_l was his y ’th percentile and gave no further points of the cdf below this, we assumed that a probability mass of y was distributed evenly across $\{c_l - 1, c_l - .5, c_l\}$. In the case of an unknown upper limit, if an expert only indicated that c_u was his y th percentile and gave no further arguments for the cdf above it, we assumed that a probability mass of $1 - y$ was distributed evenly across $\{c_u + .5, c_u + 1\}$.

Following this procedure we arrived at probability distribution functions for all 14 experts conditional on all three Scenarios. The implied cumulative distribution functions are depicted in Figure 3.

A3: Tables and Figures not in main text

2030 solar-technology cost c (2005 USc\$/kWh)	Benefit $B(c)$ (US\$ 10 ⁹)
2	189.90
2.5	170.76
3	151.26
3.5	131.74
4	112.12
4.5	92.29
5	71.47
5.5	50.64
6	29.27
6.5	23.59
7	12.32
7.5	3.67
8	1.76
> 8	0

Table 1: EU discounted consumption improvement as a function of 2030 solar-power cost

R&D Scenario r	Opportunity Cost $O(r)$ (US\$ 10 ⁹)
r_1	3.67
r_2	5.51
r_3	7.35

Table 2: Discounted opportunity cost of R&D Scenarios

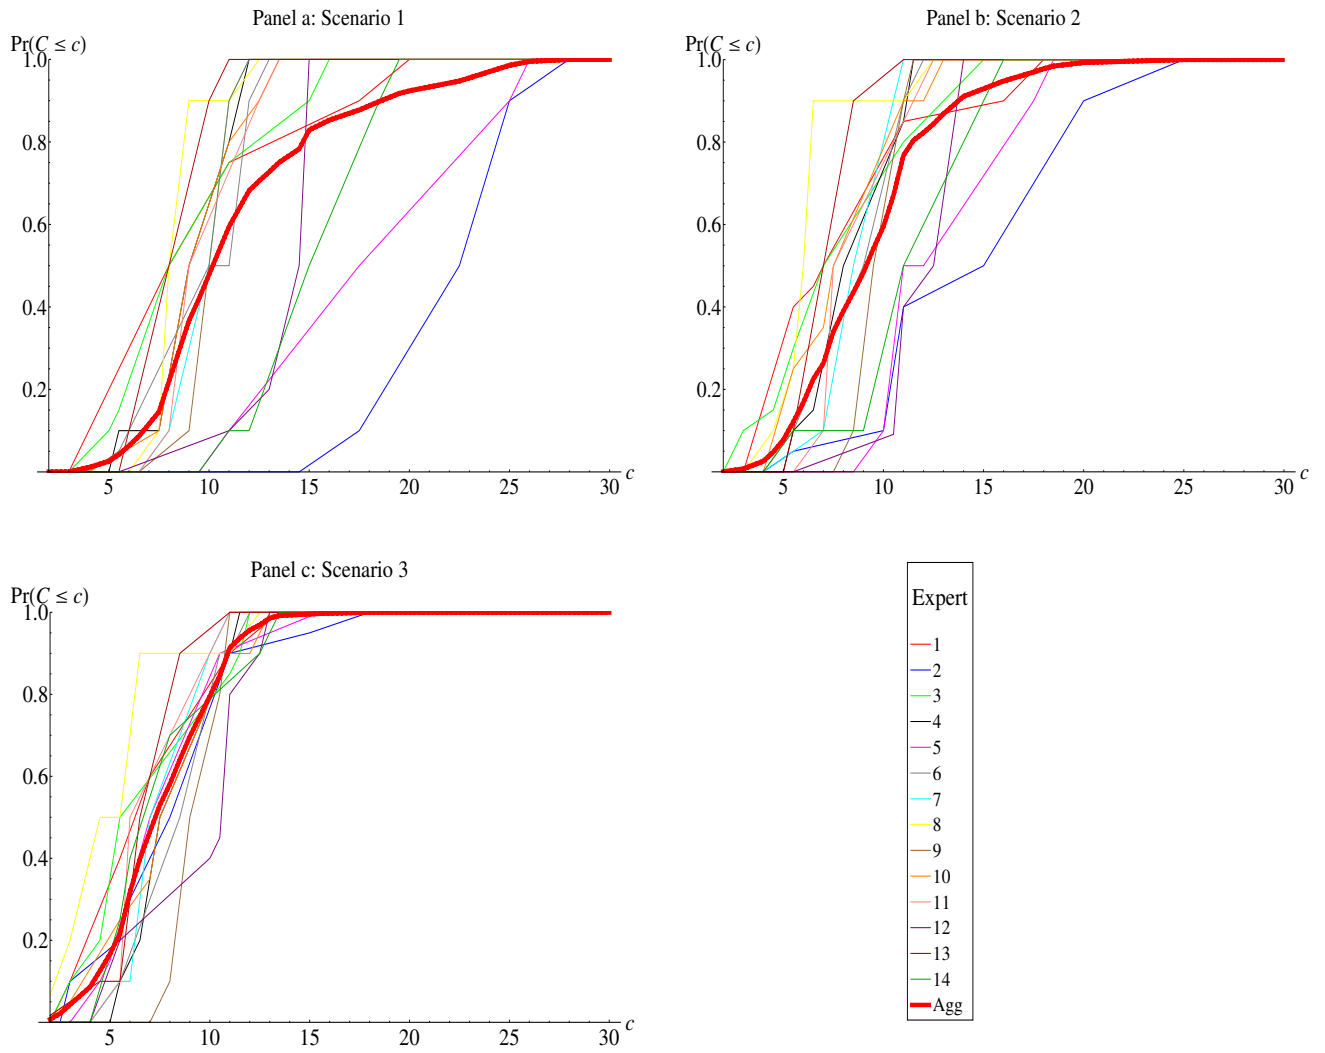


Figure 2: Expert and aggregate cdfs of the 2030 cost of solar technology conditional on the three R&D Scenarios. Recall that the cdf's domain is $\{2, 2.5, \dots, 29, 29.5, 30\}$. Cost is measured in 2005 US\$/kWh.




Figure 3: Worst and Best-Case net payoffs (benefits minus opportunity cost) for the three R&D scenarios. Net payoffs are measured in US\$10⁹.

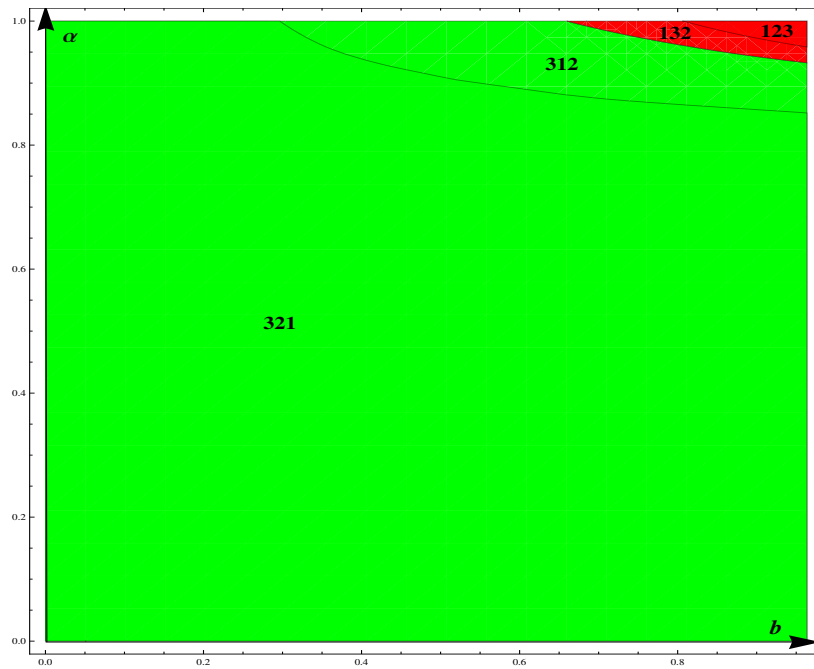


Figure 4: Comparison of the three R&D scenarios over all values of b and α .

References

- [1] D. Ahn (2008), “Ambiguity without a State Space,” *Review of Economic Studies*, 75, 3–28.
- [2] K. Arrow and L. Hurwicz (1972), “An Optimality Criterion for Decision-Making under Ignorance,” in C.F. Carter and L.J. Ford (eds.) *Uncertainty and Expectations in Economics: Essays in Honour of G.L.S. Shackle* (Oxford: Basil Blackwell), 1-11.
- [3] T. Asano (2010). Precautionary principle and the optimal timing of environmental policy under ambiguity. *Environmental and Resource Economics*, 47, 173-196.
- [4] S. Athanassoglou, and A. Xepapadeas (2012), “Pollution control with uncertain stock dynamics: When, and how, to be precautionary,” *Journal of Environmental Economics and Management*, 63, 304-320.
- [5] A. Baillon, L. Cabantous, and P. Wakker (2012), “Aggregating imprecise or conflicting Beliefs: An experimental investigation using modern ambiguity theories,” *Journal of Risk and Uncertainty*, 44(2), p.115-147
- [6] E. Baker, H. Chon, J. Keisler (2009), “Advanced solar R&D: Combining economic analysis with expert elicitations to inform climate policy,” *Energy Economics*, 31, S37-S49.
- [7] E. Baker, H. Chon, J. Keisler (2009), “Carbon capture and storage: combining economic analysis with expert elicitations to inform climate policy,” *Climatic Change*, 96, 379-408.
- [8] E. Baker and J. Keisler (2011), “Cellulosic biofuels: Expert views on prospects for advancement,” *Energy*, 36, 595-605.
- [9] V. Bosetti, C. Carraro, M. Galeotti, E. Massetti and M. Tavoni (2006), “WITCH: A World Induced Technical Change Hybrid Model,” *Energy Journal*, 2, 13-38.
- [10] V. Bosetti, M. Catenacci, G. Fiorese, and E. Verdolini (2012). “The future prospect of PV and CSP solar technologies: An expert elicitation survey,” *Energy Policy*, 49, 308-319.
- [11] R. Clemen and R. Winkler (1999), “Combining Probability Distributions from Experts in Risk Analysis,” *Risk Analysis*, 19, 187–203.
- [12] R. Clemen and R. Winkler (2007), “Aggregating probability distributions,” In: W. Edwards, R. Miles, & D. von Winterfeldt (Eds.) *Advances in Decision Analysis* .(pp. 154-176). Cambridge , UK : Cambridge University Press.
- [13] R. Cooke (1991), *Experts in Uncertainty: Opinion and Subjective Probability in Science*, Oxford University Press, USA.
- [14] A.E. Curtright, G. Morgan, D.W. Keith (2008), “Expert Assessments of Future Photovoltaic Technologies,” *Environmental Science and Technology*, 42, 9031-9038.
- [15] J. Eichberger, S. Grant, D. Kelsey, and G. Koshevoye (2011), “The α -MEU model: A comment,” *Journal of Economic Theory*, 146, 1684–1698.
- [16] G. Fiorese, M. Catenacci, E. Verdolini, and V. Bosetti (2013), “Advanced biofuels: Future perspectives from an expert elicitation survey,” *Energy Policy*, 56, 293–311.

- [17] Gajdos, T., Hayashi, T., Tallon, J. M., and Vergnaud, J. C. (2008), "Attitude toward imprecise information" *Journal of Economic Theory*, 140, 27-65.
- [18] Gajdos, T, and J. C. Vergnaud (2013). "Decisions with conflicting and imprecise information." *Social Choice and Welfare* 41, 427-452.
- [19] P. Ghirardato (2001), "Coping with Ignorance: Unforeseen Contingencies and Non-Additive Uncertainty," *Economic Theory*, 17, 247–276.
- [20] P. Ghirardato, F. Maccheroni, and M. Marinacci (2004), "Differentiating Ambiguity and Ambiguity Attitude," *Journal of Economic Theory*, 118, 133-173.
- [21] I. Gilboa, and M. Marinacci (2012), "Ambiguity and the Bayesian Paradigm," *Advances in Economics and Econometrics: Theory and Applications*, Tenth World Congress of the Econometric Society.
- [22] I. Gilboa and D. Schmeidler (1989), "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141–153.
- [23] J.-Y. Jaffray (1989), "Linear Utility Theory for Belief Functions," *Operations Research Letters*, 8, 107–112.
- [24] P. Klibanoff, M. Marinacci, and S. Mukerji (2005), "A Smooth Model of Decision Making Under Ambiguity," *Econometrica*, 73, 1849–1892.
- [25] R. Knutti, R. Furrer, C. Tebaldi, J. Cermak, and G. A. Meehl (2010). "Challenges in combining projections from multiple climate models." *Journal of Climate* 23, 2739-2758.
- [26] Kriegler, E., Hall, J. W., Held, H., Dawson, R., Schellnhuber, H. J. (2009). Imprecise probability assessment of tipping points in the climate system. *Proceedings of the National Academy of Sciences*, 106(13), 5041-5046.
- [27] D. Ludwig, R. Hilborn, and C. Walters (1993). "Uncertainty, resource exploitation, and conservation: lessons from history." *Science*, 260, 17-18.
- [28] Meinshausen, Malte, Nicolai Meinshausen, William Hare, Sarah CB Raper, Katja Frieler, Reto Knutti, David J. Frame, and Myles R. Allen. "Greenhouse-gas emission targets for limiting global warming to 2 C." *Nature* 458, no. 7242 (2009): 1158-1162.
- [29] Millner, A., Calel, R., Stainforth, D. A., MacKerron, G. (2013). Do probabilistic expert elicitations capture scientists uncertainty about climate change?. *Climatic change*, 116(2), 427-436.
- [30] A. Millner, S. Dietz, and G. Heal (2013). Scientific ambiguity and climate policy. *Environmental and Resource Economics*, 55, 21–46.
- [31] M.G. Morgan, and M. Henrion (1990), *Uncertainty: A guide to dealing with uncertainty in quantitative risk and policy analysis*, Cambridge University Press, New York.
- [32] M.G. Morgan, I. Pitelka, and E. Shevliakova (2001), "Elicitation of Expert Judgments of Climate Change Impacts on Forest Ecosystems," *Climatic Change* 49, 279-307.
- [33] M. G. Morgan, P. J. Adams, and D.W. Keith (2006), "Elicitation of Expert Judgments of Aerosol Forcing," *Climatic Change*, 75, 195-214

- [34] W. Nordhaus (1994), “Expert Opinion on Climatic Change,” *American Scientist* 82, 45-51.
- [35] W. Nordhaus. *A question of balance: Weighing the options on global warming policies*. Yale University Press, 2008.
- [36] A. O’Hagan, C.E. Buck, A. Daneshkhan, J.R. Eiser, P.H. Garthwaite, D.J. Jenkinson, J.E. Oakey, and T. Rakow, (2006), *Uncertain judgments: Eliciting experts probabilities*, John Wiley and Sons, Ltd., Chichester.
- [37] W. Olszewski (2007), “Preferences Over Sets of Lotteries,” *Review of Economic Studies*, 74, 567-595.
- [38] D. Popp (2004), “ENTICE: Endogenous Technological Change in the DICE Model of Global Warming,” *Journal of Environmental Economics and Management*, 48, 742–768.
- [39] M. Rostek (2010), “Quantile Maximization in Decision Theory,” *Review of Economic Studies*, 77, 339–371.
- [40] L. J. Savage (1954), *The Foundations of Statistics*, Wiley, New York.
- [41] N. Treich (2010). The value of a statistical life under ambiguity aversion. *Journal of Environmental Economics and Management*, 59, 15-26.
- [42] A. Tversky and D. Kahneman (1974), “Judgment under Uncertainty: Heuristics and biases,” *Science*, 185, 1124–1131.
- [43] A. Tversky and D. Kahneman (1983), “Extensional vs. Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment,” *Psychological Review* 90, 293–315.
- [44] U.S. EPA (2011), Expert Elicitation Task Force White Paper. available at <http://www.epa.gov/stpc/pdfs/ee-white-paper-final.pdf>
- [45] G. Vardas and A. Xepapadeas (2010), “Model Uncertainty, Ambiguity and the Precautionary Principle: Implications for Biodiversity Management,” *Environmental and Resource Economics*, 45, 379–404.
- [46] D. Vaughan and J. Spouge (2002), “Risk Estimation of Collapse of the West Antarctic Ice Sheet,” *Climatic Change* 52, 65–91.
- [47] R.T. Woodward and R.C. Bishop (1997), “How to Decide When Experts Disagree: Uncertainty-Based Choice Rules in Environmental Policy,” *Land Economics*, 73, 492-507.
- [48] Yaari, M. E. (1969). Some remarks on measures of risk aversion and on their uses. *Journal of Economic theory*, 1(3), 315-329.
- [49] K. Zickfeld, A. Levermann, M.G. Morgan, T. Kuhlbrodt, S. Rahmstorf, and D.W. Keith, (2007), “Expert judgements on the response of the Atlantic meridional overturning circulation to climate change,” *Climatic Change*, 82, 235-265.
- [50] K. Zickfeld, M.G. Morgan, D. Frame, D. Keith (2010), “Expert Judgments about Transient Climate Response to Alternative Future Trajectories of Radiative Forcing,” *Proceedings of the National Academy of Science* 107, 12451-12456.