

Finding common ground when experts disagree: Belief dominance over portfolios of alternatives

by

Erin Baker, Valentina Bosetti, Ahti Salo

Abstract

We address the problem of choosing a portfolio of policies under “deep uncertainty.” We introduce the idea of *belief dominance* as a way to derive a set of non-dominated portfolios and robust individual alternatives. Our approach departs from the tradition of providing a single recommended portfolio; rather, it derives a group of good portfolios. The *belief dominance* concept allows us to synthesize multiple expert- or model- based beliefs by uncovering the range of alternatives that are intelligent responses to the range of beliefs. This goes beyond solutions that are optimal for any specific set of beliefs to uncover other defensible solutions that may not otherwise be revealed. We illustrate our approach using an important problem in the climate change and energy policy context: choosing among clean energy technology R&D portfolios. We demonstrate how the belief dominance concept can reveal portfolios and alternatives that would otherwise remain uncovered.

I. Introduction

An important reason for why governments have been slow to address climate change is the uncertainty that surrounds it. Some groups have exploited this uncertainty to impede forward movement to address climate change as a global problem (Oreskes & Conway, 2011). In addition, both ends of the political spectrum have spent considerable time and resources arguing against specific solutions, with, for example, some on the right opposing solar and wind energy, and some on the left opposing nuclear and carbon capture. These arguments have led to a conservative approach, with few solutions moving forward at a speed that is needed to avoid serious climate damages (IPCC, 2014).

In this paper, we address the problem of “deep uncertainty”. This problem—which is pervasive in climate change and several other instances of collective decision making—refers to a situation in which there is significant disagreement about probability distributions over relevant outcomes (McInerney et al 2012). A frequently discussed example is that of multiple beliefs regarding the climate response to CO₂ atmospheric concentrations, as measured by climate sensitivity (Caldeira et al, 2003). On another

front, Tol (2013) and Drouet et al (2015) provide examples of multiple beliefs over the socio-economic damages associated to changes in the climate.

In order to ground our paper in an example, we present a climate-related proof-of-concept: energy technology R&D portfolios in response to climate change. There are multiple beliefs over the future performance of key mitigation technologies, which, in turn, can be mapped into beliefs over the overall cost of climate mitigation or clean energy policies. Multiple studies report different distributions over the future costs of solar PV, nuclear, biofuels, etc., often conditional on specific policy interventions. The problem decision-makers face is to use these multiple views –which are often in disagreement–to define an optimal (or at least good) portfolio for pursuing energy-related research and development. In principle, policy makers would want to know the best composition of the energy innovation portfolio to meet their objective, be it reducing energy imports or Greenhouse gas emissions from the energy sector.

The broad question we tackle is how to approach deep uncertainty in the development of public policy strategies, where deep uncertainty is defined as a situation in which experts or models generate conflicting beliefs over future states of the world. The most traditional approach (particularly with regard to experts) is to aggregate expert judgements into a single distribution. There is a large literature on this topic. In a comprehensive review, Clemen & Winkler (1999) discuss both behavioral and mathematical aggregation, and conclude that while “no single process is best in all circumstances,” mathematical aggregation methods appear to have a “slight edge” over behavioral ones, and the simple average performs quite well. Cooke and Goossens (2008) have shown that weighting experts based on the results of test questions can increase calibration. Recent work has suggested that other forms of mathematical aggregation, such as using medians (Hora et al 2013) or averaging quantiles (Lichtendahl et al 2013), have attractive properties. This literature is generally agnostic about the decision context – once aggregated, the beliefs are portable from one context to another. These aggregated beliefs can then be used with traditional approaches of decision making under uncertainty (e.g. Baker and Solak 2014, Kelly and Kolstad 1999, Keller et al 2004). These approaches, however, have been criticized for providing solutions that appear too certain or are lacking in “external consistency” (Milner et al 2014;

Gilboa et al 2009). That is, they generally provide a mathematical solution to the disagreement among experts; and result in a single best solution.¹

A second approach when there is deep uncertainty is to communicate to decision makers the full range of judgements and model results (as for example Morgan and Herion, 1990). This approach is cognizant of the decision makers in a way that differs from the aggregation methods above, but it generally provides elicited distributions and/or model results that are portable from one context to another. Moreover, this approach leaves open the problem of what to do with the multiple beliefs: how the decision makers can use them to actually inform decision-making.

A third method builds on the full set of multiple beliefs, but integrates the decision-maker preferences about missing or ambiguous information directly into the decision problem. This is done by means of non-traditional decision rules, for example by applying the machinery of ambiguity aversion in economics (see Milner et al 2013 or Berger et al., 2015 for applications to the climate change policy context) or Robust Optimization in operations research (See Gabrel et al 2014 and Bertsimas et al 2010 for recent reviews). These methods can best be understood as worst-case analyses: they strive to find a solution that performs well across a wide range of parameter values or beliefs². This approach does not explicitly synthesize the beliefs separately from the problem (like the other two above which can be portable from one context to another). Rather, the beliefs are synthesized within the context of the problem itself and lead to a specific recommendation for action. The non-expected utility decision criteria used in this approach, however, have been criticized for not being internally consistent (Al Najjar and Weinstein 2009, Baker and Regnier 2015).

Both the traditional approach and the ambiguity aversion/robust optimization approach provide a mathematical solution to the disagreement among experts; and result in a single best decision recommendation. We argue that in cases where there is disagreement (arising, for instance, from multiple plausible perspectives concerning relevant beliefs), there are benefits to adopting an approach in which analysts do not necessarily provide a single “best” decision recommendation (like in the first

¹ While behavioral aggregation, such as consensus or the Delphi approach, do not rely on a strictly mathematical aggregation, they do provide a single aggregation and have been criticized as being sensitive social pressures (Clemen and Winkler 1999).

² This is a simplification of course. See Ben-Tal et al 2010 for a nice example of a paper that allows for smooth tradeoffs between optimization and worst case avoidance.

and third approaches), but go beyond simply communicating the range of disagreement (like the second approach).

This general idea – of providing and evaluating multiple alternatives rather than a single best decision – has been applied in a set of bottom-up exploratory approaches such as Robust Decision Making (RDM) (Rosenhead et al 1972, Lempert and Collins 2007), Decision Scaling (Brown et al 2012), and Info Gap (Ben-Haim 2004); see Kalra et al (2014) for a discussion of how these types of models can help lead to agreement over decisions. These methods typically analyze a small set of pre-defined alternatives for robustness and then suggest possible new alternatives based on the analysis (see Herman et al 2015 for a review). These approaches synthesize the range of beliefs and models within a decision context by visually communicating the range of possible outcomes implied by the range of beliefs.

The framework we describe in this paper, which we call Robust Portfolio Decision Analysis, is complementary to this last set of approaches. We introduce a method of synthesizing beliefs that is integrated with the decision context and provides a set of defensible alternatives, which can be further analyzed with the bottom-up exploratory approaches. In this sense, it departs from the tradition of providing a single recommended solution; rather, it derives a group of good alternatives. In order to identify these alternatives, we introduce a new dominance concept – *belief dominance* – that allows us to synthesize beliefs by uncovering the range of alternatives that are intelligent responses to the range of beliefs. The use of this dominance concept is more powerful than sensitivity analysis in that it will always include the optimal solution for any relevant beliefs (see Wallace 2000 for a discussion of sensitivity analysis); moreover, it goes beyond the solutions which are optimal for any specific set of beliefs to uncover other defensible solutions that may not otherwise be revealed.

Providing a set of good alternatives rather than a single optimal solution allows decision makers to consider a wider range of information, which may be non-quantitative or incomplete. Such information—which may pertain to stakeholders’ preferences, developments in technological progress, or the consequences of technological choices—can then be incorporated explicitly or implicitly in the decision making process. Indeed, while modelers believe their *models are useful*, most decision makers are well aware that *all models are wrong* (Box and Draper, 1987), if only because there are many aspects of the real world that cannot be modeled. This recognition suggests that there is good reason to use models to gain insights and to build a set of good solutions, helping decision makers to make choices while leaving

room for flexibility to apply a wider range of criteria. In summary, the analysis and generation of several non-dominated alternatives from which decision makers can choose characterizes our approach.

Another key distinction of our approach is that we focus on portfolios of individual alternatives. These other methods are agnostic about the specific alternatives but in practice tend to focus on individual alternatives. The focus on portfolios has a significant advantage in situations where there are multiple stakeholders, such as the world of public policy. Our method can highlight individual alternatives that are robust in the sense that they are part of the portfolio regardless of the individual beliefs. This allows conflicting stakeholders with conflicting sets of beliefs to find some common ground, which is well known to improve the outcomes of negotiation and deliberation (Mansbridge et al 2013). When applied to climate change, this portfolio approach, which might include a varied set of mitigation and adaptation strategies, may open up the dialog to a wider group of constituencies, laying hope for a societal solution to this global challenge (Center for research on environmental decisions, 2014). Indeed, scholars in the field of public engagement have suggested that discussion focused on a broad selection of solutions may appeal to, and mobilize, a wider range of stakeholders than a sole focus on the consequences of climate change (Roser-Renouf et al, 2014).

We demonstrate our approach by applying it to the problem of choosing publically-funded energy technology R&D portfolios using data coming from the TEaM project, which collected multiple experts' beliefs on the climate-energy technology R&D domain (Baker et al 2015).

In the next section, we define the theoretical framework and draw a parallel with stochastic dominance and multi-objective decision making. Section III introduces a specific application of the methodology to the case of energy R&D portfolio selection. Section IV discusses the flexibility and extensions of the presented approach, and Section V concludes.

II. Robust Portfolio Decision Analysis – theoretical framework

Let us start from the broadest definition of the problem: identifying individual alternatives among portfolios that are robust to a range of beliefs about the outcomes associated to each alternative. There are two pieces to the theoretical framework that we introduce here. The first is the concept of belief

dominance, defined so that a portfolio A dominates B if A is preferred to B for all probability distributions that represent plausible beliefs concerning the outcomes of these alternatives. In our example, this equates to finding portfolios of R&D investments in energy technologies that are preferred to other portfolios across the full set of experts' beliefs concerning the performance of R&D on the future energy technologies costs. From this information, we build the set of non-dominated portfolios. The second piece of the framework allows us to move from the set of non-dominated portfolios to derive implications about individual strategies composing the portfolio. Again, to use our specific example, this represents shifting the focus to individual R&D investments decisions to find, for example, those that are present in all non-dominated portfolios or those that are never present.

II.1 Belief Dominance

Consider the following generic decision model

$$\max_{\mathbf{x}} \int U(\mathbf{x}; z) f(z; \mathbf{x}) dz \quad (0)$$

where \mathbf{x} is a vector of decision variables, z is a realization of the random variable Z with probability distribution f (that may or may not depend on the decision variables \mathbf{x}), and U is an objective function (which may or may not include risk aversion). Belief dominance compares alternatives over sets of beliefs.

We define *belief dominance* as follows: an alternative \mathbf{x} *belief dominates* alternative \mathbf{x}' over a set Φ of probability distributions if and only if

$$\int U(\mathbf{x}; z) f(z; \mathbf{x}) dz \geq \int U(\mathbf{x}'; z) f(z; \mathbf{x}') \quad \forall f \in \Phi \quad (0)$$

and the inequality is strict for at least one f . Thus, this definition is specific to the decision problem as defined by U , which represents the mapping of the primitives (decisions and random variables) to metrics of interest (such as the cost of achieving a climate target) and includes the decision maker's preferences, most notably his or her attitudes towards risk. Note that we explicitly allow for the possibility that the probability distribution is contingent on the decision \mathbf{x} .

For intuition, we put this in context with other common dominance concepts:

- *Belief*: alternative x dominates alternative x'

$$\int U(\mathbf{x}; z) f(z; \mathbf{x}) dz \geq \int U(\mathbf{x}'; z) f(z; \mathbf{x}') \quad \forall f \in \Phi \quad (1)$$

- *Stochastic*: distribution f dominates distribution g

$$\int U(\mathbf{x}; z) f(z) dz \geq \int U(\mathbf{x}; z) g(z) \quad \forall U \in V \quad (2)$$

where V is a set of utility functions representing a type of risk preference, such as “concave functions” representing risk aversion;

- *Pareto*; alternative x dominates alternative x'

$$\int U_i(\mathbf{x}; z) f(z) dz \geq \int U_i(\mathbf{x}'; z) f(z) \quad \forall U_i \quad (3)$$

where U_i represent multiple objective functions.

Stochastic dominance compares probability distributions over sets of objective functions. If the set V in inequality (3) is the set of functions increasing in z , then the central inequality defines first order stochastic dominance; if V is the set of functions increasing and concave in z , this is second order stochastic dominance; if V is the set of concave functions (with positive domain), this defines an increase in risk in the Rothschild-Stiglitz sense.

Pareto dominance, represented in inequality (5), compares alternatives over sets of objectives. Often these objectives are discrete criteria, such as “cost”, “safety”, “reliability”; in this case such problems are often referred to as Multi-Criteria Decision Making or MCDM. For consistency, we have written the Pareto Dominance inequality in terms of expected values; often Pareto Dominance is used in a deterministic setting, in which $f(z)$ would simply put all weight on one, deterministic point.

Thus, belief dominance is similar to Pareto dominance in that it compares alternatives. It is similar to stochastic dominance in that it refers to probability distributions. It differs from both definitions in that it refers to dominance over beliefs rather than over preferences. That is, in our context, the disagreement or uncertainty is not over preferences, but over beliefs about the outcomes of alternatives, represented by sets of probability distributions. When there are multiple different beliefs (for instance, when they are stated by experts who do not agree), we suggest that analysis should yield a set of non-dominated alternatives just as in the cases of stochastic and Pareto dominance; and

furthermore, that broader disagreement over beliefs should lead to a larger set of non-dominated alternatives.

Again parallel to the other dominance concepts, we are interested in applying belief dominance to determine the set of decisions \mathbf{x} which are non-dominated; that is, all those decisions for which there is no other decision \mathbf{x}' that would be preferred to \mathbf{x} across the full range of beliefs expressed by experts. Methodologically, this has parallels to the identification of Pareto efficient alternatives when considering multiple objectives (we note that Keeney (1996) defines “objectives” as evaluation criteria, which include a direction of preference, such as minimizing costs or maximizing health benefits). By way of analogy, the experts’ beliefs can be viewed in the role of objectives, and hence an alternative \mathbf{x} is non-dominated if there is no other alternative \mathbf{x}' such that the expected objective function of \mathbf{x}' is higher than that of \mathbf{x} for all beliefs. Conversely, if alternative \mathbf{x} is dominated by \mathbf{x}' , then the expected objective function of \mathbf{x}' is higher than that of \mathbf{x} under the beliefs of all experts (i.e., in their individual evaluations of these alternatives, all experts would conclude that \mathbf{x}' is either strictly better than \mathbf{x} , or at least as good as \mathbf{x} in terms of the objective U).

In our analysis, we determine the set of alternatives that are non-dominated across all individual beliefs. Here we provide a theorem showing that, if an alternative is dominated over a finite set of beliefs, it is also dominated over the convex combination of these beliefs.

Theorem. Let f_1, f_2, \dots, f_n be a finite set of beliefs, represented as probability distributions, and let Φ be the set of convex combinations of the beliefs such that $\Phi = \{f = \lambda_1 f_1 + \lambda_2 f_2 + \dots + \lambda_n f_n \mid \text{for all } \lambda_i, i=1, \dots, n \text{ such that } \lambda_i \geq 0, \sum_i \lambda_i = 1\}$. Alternative \mathbf{x} belief dominates \mathbf{x}' for all discrete beliefs $i=1, \dots, n$, if and only if \mathbf{x} belief dominates \mathbf{x}' over Φ .

This theorem—which is formally proved in the Appendix —implies that if the presence of dominance is established for all individual beliefs, then dominance also holds for all *combinations* of such beliefs. Conversely, if dominance between \mathbf{x} and \mathbf{x}' does not hold in the belief set Φ , defined as the convex set of a finite number of experts’ beliefs, there must exist at least two distinct beliefs which rank the alternatives \mathbf{x} and \mathbf{x}' differently.

This theorem illustrates the power of this method with respect to traditional parametric sensitivity analysis. It has long been understood that sensitivity analysis – finding the optimal solution under a

number of candidate probability distributions – is in no way guaranteed to reveal the actual optimal solution (See Wallace (2000) for seminal paper). That is, the optimal solution is not guaranteed to be contained in the space spanned by the deterministic solutions (or the solutions of individual probability distributions). This theorem indicates that the non-dominated set does not have this problem: the optimal solution for any convex combination of the candidate distributions is guaranteed to be in the non-dominated set. Any solution that is optimal for any probability distribution that is a convex combination of the candidate distributions will be part of the non-dominated set.

II.2 Deriving recommendation for alternatives from portfolio-level analyses

Our approach builds on the ideas of Robust Portfolio Modeling (RPM; Liesiö et al., 2007, 2008) which supports the selection of a portfolio of alternatives (such as R&D projects) from a large set of candidates. Specifically, its extension to scenario analysis (Liesiö and Salo, 2012) employs set inclusion to capture uncertainties about the decision maker’s risk preferences and beliefs by accommodating (1) sets of feasible utility functions over outcomes and (2) sets of feasible probability distributions over distinct scenarios. Results are obtained by determining which portfolios are non-dominated, in the sense that there does not exist any other portfolio that would be at least as good for all feasible combinations of utility functions and probabilities, and strictly better for some such combination.

The conceptual breakthrough in RPM is to analyze the set of non-dominated portfolios to inform choices among individual alternatives by dividing these alternatives into three categories. First, those alternatives that are contained in all non-dominated portfolios belong to the *core*. Second, alternatives that are not contained in any non-dominated portfolios are *exterior*. Finally, the *borderline* consists of alternatives that are included in some but not all non-dominated portfolios. To define this mathematically, let projects be indexed by $i = 1..I$, and define $x_i = 1$ if project i is invested in and 0 otherwise. Let ND denote the set of non-dominated portfolios. Then we can define the three sets as follows:

$$\begin{aligned}
 core &\equiv \{i \mid x_i = 1 \ \forall \vec{x} \in ND\} \\
 ext &\equiv \{i \mid x_i = 0 \ \forall \vec{x} \in ND\} \\
 bord &\equiv \{i \mid i \notin core \text{ and } i \notin ext\}
 \end{aligned}
 \tag{4}$$

Individual Projects:	a	b	c	d	e	f
Portfolios						
1	1	0	0	1	1	0
2	1	0	1	1	1	0
3	1	0	0	1	1	0
4	0	0	1	1	0	1
5	0	0	0	1	0	1
6	0	0	1	1	0	1

Table 1: illustration of the composition of the six non-dominated portfolios composed of individual projects a, b,...,f.

Table 1 provides an illustrative example, in which the 6 rows represent the 6 non-dominated portfolios; and the projects *a-f* can be invested in or not. In this case project *b* is in the exterior, project *d* is in the core; all other projects are in the borderline.

An important theoretical result is that when uncertainties are reduced—in the sense that the sets of feasible probability distributions become smaller—all core and exterior alternatives stay in their respective sets (see Theorem 2 in Liesiö and Salo, 2012). As a result, recommendations concerning the selection of core alternatives and the rejection of exterior alternatives are robust, because these recommendations would stay valid even if additional information is obtained. For example, an individual technology investment that is in the core over a finite set of probability distributions will remain in the core for combinations of feasible probability distributions, including any subset of these distributions. Thus, research aimed at deriving recommendations that are more conclusive should be focused on the borderline alternatives: for instance, it is possible to analyze if these borderline alternatives can be enhanced to make them equally attractive as some core alternatives (Gregory and Keeney, 1994); or if gathering more information about the borderline alternatives allows them to be moved into the core or the exterior.

On the other hand, this result implies that when additional perspectives are added, making the feasible set of probability distributions larger, the core and exterior sets may become smaller. In the extreme case, all alternatives will belong to the borderline, which makes it harder to differentiate which alternatives are better than others are. We discuss this further in the next section.

II.3 Refinements and Robustness Concepts

The intention of our framework is to identify common ground in order to catalyze discussion about near term actions. The core and exterior sets represent this common ground. The borderline set represents the individual alternatives over which negotiation needs to take place. If the borderline set is too large, then this framework may not provide enough common ground to get started. Thus, in this section we briefly discuss some potential refinements to this method, which can be used to catalyze discussion on how to reduce the size of the borderline set. The refinements we discuss are in three categories: (1) imposing constraints on the set of relevant beliefs; (2) eliminating some Non-Dominated portfolios to create a more robust set; and (3) employing measures about the individual alternatives in the context of portfolio selection.

Constraining Beliefs. Decision makers may want to impose constraints on the set of relevant beliefs. This can be done by associating weights over the sources of information (e.g., experts or models) from which the probability distributions have been derived, and then imposing constraints on these using different rules. Any individual rule would result in a single aggregated distribution; a set of acceptable rules would result in a set of aggregated distributions. For example, the set of rules might include equal weight on all priors, weights proportional to the number of experts represented by each prior; and weights depending on the subjective quality of each prior (see e.g. Vilkkumaa et al., 2015). For another example, weights over experts could be constrained so that no single expert will have more than 50% of the total weight. In this case, the set Φ that spans all the relevant beliefs will be $\{f = \lambda_1 f_1 + \lambda_2 f_2 + \dots + \lambda_n f_n \mid 0.50 \geq \lambda_i \geq 0, \sum_i \lambda_i = 1\}$. The weight constraint $0.50 \geq \lambda_i$ reduces the set of probability distributions, which will typically make the sets of core and exterior alternatives larger.

Robustness Concepts. Another approach is to combine this framework with concepts of robustness from the literature. The results of this framework can be visually presented to decision-makers in a way that allows them to iteratively remove less robust portfolios from the set of non-dominated portfolios. According to Lempert et al 2006 “A robust strategy performs relatively well—compared to alternatives—across a wide range of plausible futures.” That is, a less robust portfolio would perform badly under some beliefs. However, there are multiple ways to define in what sense a strategy performs relatively well or relatively poorly. In the spirit of this framework, we suggest that analysts present the data in multiple ways in order to support and induce discussion. In particular, we suggest

presenting (1) the expected objective values of each of the portfolios under each belief; and (2) the percentage decrease in value (or increase in cost) for each belief with respect to the lowest value for each belief. Less robust portfolios would perform poorly, either in an absolute or relative way, in either or both of the presentations. We discuss this further in Section III in the context of our proof-of-concept, and tie these two concepts to common definitions of robustness in the literature.

Measures of individual alternatives. The core and exterior sets are extreme, in that they represent individual alternatives that are present or missing in every non-dominated portfolio. Another approach is to consider measures of individual alternatives that are on a continuum rather than black or white. One such measure is the Core Index (CI), which is defined as the ratio between the number of portfolios which contain an individual alternative versus the total number of non-dominated portfolios (Liesiö et al, 2007). The resulting CI values can then be employed to obtain tentative guidance as to which alternatives are most important to analyze further. For example, referring to the illustrative example presented in Table 1, project *a* has a CI of 0.5.

II.3 Comparison with Other Approaches.

How does this approach differ from others in the literature? The key aspects of this approach are that we use *theoretically sound decision criteria*, but *multiple representations of uncertainty* to screen viable *sets of alternatives*. This approach differs from the literature on ambiguity aversion and robust optimization in that those methods use some version of a worst-case analysis. Worst-case analyses can range from straightforward Maxmin rules (Gilboa & Schmeidler 1989; Ribas et al 2010), to more sophisticated methods such as Klibanoff's smooth ambiguity (Klibanoff et al 2005) or Ben-Tal et al 2010 Soft Robustness. A weakness of these methods is that recommendations based on these types of decision rules are not consistent through time. For example, Al-Najjar et al (2009) show that employing ambiguity aversion leads to results such as sensitivity to sunk costs or aversion to information³. Our approach differs from the traditional literature on dynamic decision making under uncertainty and learning (for example Baker & Solak 2014; Kelly & Kolstad 1999; Webster et al 2008) in that we explicitly recognize and accommodate disagreements over beliefs. It differs from both the above strands of literature in providing sets of alternatives for decision makers to choose among. Policy makers, in particular, are likely to find this feature very attractive, because it allows them to balance a number of

³ Siniscalchi (2009) argues that the result on sunk costs is technical rather than substantial.

criteria, including some they may not be able to quantify. On the other hand, the framework as presented here is a simple one with limited recourse: we do not explicitly model or allow for later revisions of key decisions (such as second stage readjustments in the allocation of R&D within a portfolio of energy technologies) when more information concerning uncertain elements become available. There are examples in the literature for each of the above methods that have been developed to include recourse, with a focus on how the availability of later decisions may affect the current decision (e.g. Chen et al 2007; Milner et al 2013). Our framework can be adapted to include recourse as well, although we leave this for future work.

Our approach is most closely related to the set of bottom-up exploratory approaches to support robust decision making, including RDM (Lempert 2002), Decision scaling (Brown et al 2012), and info-gap (Ben-Heim 2004). See Herman et al. 2015 for a review. All of these approaches evaluate a set of alternatives against robustness criteria using scenario discovery to identify key uncertainties. They do not generally provide a single best solution, but rather focus on clearly communicating how the implications of different alternatives compare in terms of robustness. Our approach complements these approaches in that we use available probabilistic information to derive a good set of alternatives that can then be analyzed with the above methods.

An approach somewhat parallel to ours is Many Objective Robust Decision Making, MORDM (Kasprzyk et al 2013; Hadka et al 2015) which, like our method, uses optimization techniques to identify a set of good alternatives for subsequent analysis with RDM. MORDM, however, as indicated by its name, focuses on cases with multiple objectives, using Pareto Optimality as its dominance criteria. In future work, our method could be combined with MORDM to produce a set of alternatives that are non-dominated in terms of both objectives and beliefs.

III. Application: Energy Technology R&D Portfolio in Response to Climate Change.

We apply our framework to the question of how to allocate research funds across a wide variety of energy technologies with varying potential for improvement and differing impacts on the economy and environment.

This is a complex research question, which has been approached through different avenues, including (i) the development of a broad range of integrated assessment models (IAMs)⁴ and (ii) multiple studies of expert judgments on the potential for technological change (Anadon, et al., 2012; Anadón, et al., 2014a; Baker & Keisler, 2011; Baker, et al., 2009b; Baker, et al., 2009a; Baker, et al., 2008; Bosetti, et al., 2012; Catenacci, et al., 2013; Chan, et al., 2011; Fiorese, et al., 2013). The IAMs have been useful for developing insights on the relative importance of technologies and the speed of their adoption (see Clarke et al. 2014 for a complete review). Nevertheless, there are considerable challenges from the viewpoint of decision and policymaking, including the large number of assumptions that are required and the significant uncertainties associated with these assumptions. Studies of expert judgments, on the other hand, have provided explicit probability distributions over the potential for technological change; but there are a number of independent and disparate studies, and thus incorporating them into the already computationally-complex IAMs becomes a challenge. In this setting, we explore how these two individual approaches can be combined in an integrative framework to derive robust model-based conclusions while recognizing the uncertainties that have been expressed by multiple stakeholders (see Figure 1 for a diagram of the decision process).

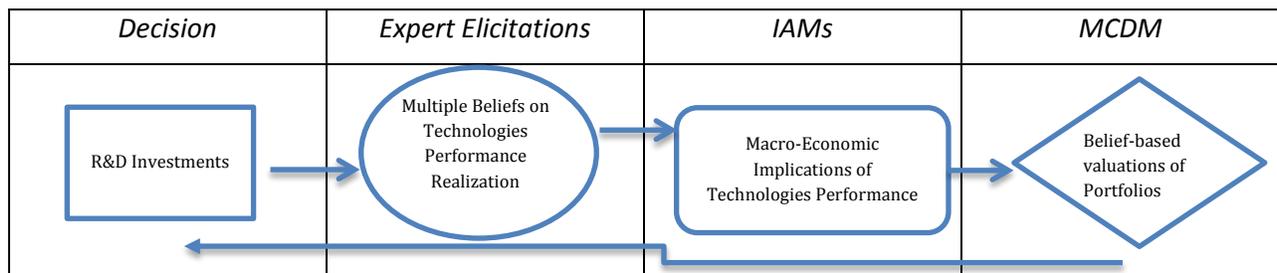


Figure 1: An influence diagram of the decision problem. Square nodes represent decisions; oval nodes uncertainties; rounded squares model calculations; and diamond nodes the objective value.

We use data on the overall welfare implications of future technologies' performance as estimated by a specific IAM, GCAM (e.g. Kim et al 2006). GCAM has been extensively used to explore the potential role

⁴ <http://www.globalchange.umd.edu/iamc/>

of emerging energy supply technologies and the greenhouse gas consequences of specific policy measures or energy technology adoption and allows studying the interactions of energy technologies with each other and with the wider economy and the environment.

We integrated this model with data derived from three large expert elicitation studies of energy technologies (summarized in Baker et al 2015). These data allow us to model multiple beliefs about key energy technologies' performances, conditional on the level of R&D investments. Because beliefs over technological performances conditional on R&D investments differ across experts, and it is not known indisputably which expert(s) may be right, the problem involves deep uncertainty. We incorporate this deep uncertainty over technological prospects by applying our concept of belief dominance in deriving sets of core, exterior, and borderline investments, and illustrate how these sets can be used to inform further research into the individual alternatives and provide insights to decision makers on near term R&D actions.

III.1 Energy technology portfolio model

For this proof of concept, the problem investigated is that of allocating R&D funds across various energy technologies in the specific context of climate change. We use the net present value of the global abatement costs as a proxy of the benefits associated with technology costs realizations, as calculated through the integrated assessment model GCAM; where abatement is defined as a reduction in emissions below a Business-as-usual baseline. We concentrate on a specific climate policy objective, which is implemented as a constraint on emissions that is compatible with a given climate stabilization scenario.

The decision problem is to choose a portfolio \mathbf{x} of investments that minimizes the expected total abatement costs (TAC) plus the opportunity cost of the portfolio itself:

$$\min_{\mathbf{x}} \int \{TAC(\mathbf{z}, s) + \kappa B(\mathbf{x})\} f_{\tau}(\mathbf{z}; \mathbf{x}) d\mathbf{z}, \quad (5)$$

where \mathbf{x} is a vector defining the portfolio; \mathbf{z} is a vector of random variables representing the realization of technological performances; s is a stabilization goal, which we set to 450 ppmeq. Total abatement costs, TAC , represent the cost of achieving stabilization goal s , given technological change realization \mathbf{z} . B is the financial cost of investing in portfolio \mathbf{x} and κ is the opportunity cost multiplier for this

investment.⁵ The probability distribution f_τ over possible outcomes depends on the selected portfolio \mathbf{x} . Because there can be multiple experts, surveys, or models describing the future probabilistic evolution of technological performance \mathbf{z} as a function of R&D investment decisions, we index f over over these beliefs, represented by τ .

A probability distribution over outcomes \mathbf{z} may depend on \mathbf{x} in complex ways. In what follows we assume that the probability of a given outcome for an individual technology depends only on the investments into that specific technology. For example, the distribution over the costs of solar depends only on the investments into solar, and not into any other technology.

TAC are nonlinear in the portfolio \mathbf{x} , and there are strong dependencies between the individual technologies in the portfolio with respect to the TAC. The TAC depends on the costs and efficiencies of the technologies, represented by the vector \mathbf{z} , and it is defined as

$$TAC(\mathbf{z}, s) = \sum_t \delta^t AC_{s,t}(\mathbf{z}), \quad (6)$$

where $AC_{s,t}(\mathbf{z})$ are the annual abatement costs (in trillions of dollars) at time t , under stabilization s and δ is the discount factor.⁶

The portfolios, \mathbf{x} , consist of investments into five key energy technologies: solar PhotoVoltaics (PV), nuclear fission, Carbon Capture and Storage (CCS), electricity from biomass (“bio-electricity”), and liquid biofuels. The cost of investment $B(\mathbf{x})$ for the portfolio \mathbf{x} is the sum of the cost of investment for each individual project. The cost of investment for each individual project is the net present value of the annual cost over 20 years using a discount rate of 3%. This data is based on the same expert elicitation protocol used to collect data on future conditional cost and efficiency distributions. Table 1 reports data on R&D cost assumptions for different levels of investments. We use an opportunity cost multiplier of $\kappa=4$.

⁵ Theory suggests that the cost to society of R&D investment may be higher than the actual dollars spent. We use a value of $\kappa=4$. See Nordhaus (2002) and Popp (2006) for details

⁶ Note that GCAM only reports values for 5-year time steps. We assume abatement costs are linear between the reported years.

Solar			Nuclear			Biofuels			Bio-electricity			CCS		
Low	Mid	High	Low	Mid	High	Low	Mid	High	Low	Mid	High	Low	Mid	High
6.7	16	132	25	77	713	5.7	15	81	5.8	12	68	21	68	673

Table 2: Annual R&D expenditures, in millions of dollars.

III.2 Multiple beliefs on Technology performance

There is deep uncertainty around the outcomes of R&D. Experts do not agree as to how R&D investment projects will impact technological change; and there are multiple beliefs over how innovation efforts will impact the future technological performance of energy technologies, $f_{\tau}(\mathbf{z}; \mathbf{x})$.

For this paper, we consider three sets of probability distributions derived from three large multi-technology expert elicitation projects carried out independently by researchers at UMass Amherst (Baker & Keisler, 2011; Baker, et al., 2009b; Baker, et al., 2009a; Baker, et al., 2008), Harvard (Anadon, et al., 2012; Anadón, et al., 2014a; Chan, et al., 2011), and FEEM (Bosetti, et al., 2012; Catenacci, et al., 2013; Fiorese, et al., 2013).⁷ We also consider the aggregation of these three (referred to as Combined – see Baker et al 2015). The Combined distribution was derived using Laplacean mixing and then smoothed using a fitted piecewise cubic distribution; therefore, it is not a simple convex combination of the other three studies. This results in four prior probability distributions over the outcomes of technological change \mathbf{z} , i.e. $\tau = 1,2,3,4$. See Figure A1 in the appendix for a visualization of the multiple distributions used in this analysis.

III.3 Calculation of non-dominated sets

There exists several ways of implementing the general problem presented in equation (5). Given the specific data we are working with, we take portfolios \mathbf{x} to be vectors of binary variables, with $x_i = 1$ if project i is invested in, and 0 otherwise. Each of the 5 technologies can be invested in at a low, medium, or high level; so each technology is associated with three mutually exclusive binary variables: exactly one decision variable associated with each technology will be equal to 1. The portfolio, given the three

⁷ For this proof of concept we consider each team as a separate belief rather than each individual expert. We did this because the individual elicitations were gathered in different ways by the different teams making the individual beliefs quite difficult to standardize as compared to the aggregated beliefs.

levels of investments into the five technologies, is a 3×5 vector of binary variables; three, mutually exclusive, levels of investment by five technologies result into $3^5 = 243$ possible portfolios.

The vector of realizations \mathbf{z} contains eight components including a cost for each of the five technologies and an efficiency for CCS, biofuels, and bio-electricity⁸. Integrated Assessment Models are quite large and computationally intensive. Thus, in order to make the set of simulations computationally feasible, we use the technique of importance sampling in a new way. Using an average of the low, mid, and high Combined distribution, we randomly draw 1000 points of the random vector \mathbf{z} ; each of the 1000 possible outcomes is represented by the non-bolded vector \vec{z} . Note that a given outcome \vec{z} is not a function of the investment vector \mathbf{x} ; only the probability of \vec{z} depends on \mathbf{x} .

Each of these points was sampled with GCAM, resulting in 1000 values of TAC. We then apply importance sampling to re-calculate the probability of each point depending on the investment portfolio \mathbf{x} . Baker et al (2015) used a set of diagnostics based on Owen (2015) and found that the samples performed in the acceptable range, with the possible exception of the biofuels and CCS efficiency parameters for the UMass and Combined distribution. See Baker et al (2015) for more details.

Thus, we have a set of technology values, $\vec{z}_l, l = 1, \dots, 1000$; and the (discrete) probability of a particular technology value realization, $f_\tau(\vec{z}_l; \mathbf{x})$, depends on the elicitation study, τ , and on the portfolio, \mathbf{x} .

We define $H(\mathbf{x}; \tau)$, the discrete version of the objective function in equation 5, given a specific set of beliefs, τ , as follows:

$$H(\mathbf{x}; \tau) \equiv \sum_{l=1}^{1000} f_\tau(\vec{z}_l; \mathbf{x}) \{TAC(\vec{z}_l, s) + \kappa B(\mathbf{x})\} \quad (7)$$

We say that a portfolio \mathbf{x} *belief dominates* \mathbf{x}' if $H(\mathbf{x}; \tau) \geq H(\mathbf{x}'; \tau)$ for all τ , with a strict inequality for at least one of the beliefs. A portfolio \mathbf{x} is non-dominated if there is no portfolio that dominates it and it is strictly better than at least one portfolio.

⁸ A complication is that the Harvard probability distributions do not distinguish between biofuels and electricity from biomass. For this initial proof of concept we assume that the investment is evenly divided between the two technologies.

As the number of portfolios is small, we explore non-dominated sets by brute-force. We first calculate the expected cost for each of the 243 portfolios, as in equation (7). We then iteratively eliminate all dominated portfolios.

III.4. Results

III.4.a Initial Results

Out of the 243 possible portfolios, ten are non-dominated across the four probability distributions.

Table 3 shows the non-dominated portfolios. They are listed in ascending order of the expected cost for the combined distribution. The first five columns provide the definition of the portfolios by showing the investment level in each technology. The last four columns show the objective value under the four different probability distributions. The objective values are color coded, with the highest cost in each column the darkest red.

Portfolios	Technologies					Objectives ENPV (cost in billions of \$2005)			
	Solar	Nuc	BF	BE	CCS	Combined	Harvard	FEEM	U Mass
1	Low	High	High	High	Mid	20736	21770	24327	15509
2	Low	Mid	High	High	Mid	20768	21654	24188	15720
3	Low	High	Mid	High	Mid	20838	21929	24525	15301
4	Mid	High	High	High	Mid	20889	21588	24345	15813
5	Low	Mid	Mid	High	Mid	20912	21806	24434	15213
6	Mid	Mid	High	High	Mid	20922	21513	24163	16162
7	High	Mid	Low	High	High	21136	21325	22747	20003
8	Mid	Mid	Mid	High	Mid	21144	21659	24379	15528
9	High	High	Low	High	High	21320	21581	22901	19324
10	Low	Mid	Mid	High	Low	21491	22671	25442	15142

Table 3 Non-dominated portfolios. Columns 2-6 report the R&D investment level for each technology, Low, Mid or High. The last 4 columns report the Expected NPV of total abatement costs plus investment cost associated to each of the portfolios under the four sets of beliefs. Higher costs are emphasized by darker red colors.

Portfolio 1 is optimal under *Combined* distribution; Portfolio 7 is optimal under both Harvard and FEEM distributions; Portfolio 10 is optimal under the UMass distribution. One benefit of this framework is that it uncovers the other 6 portfolios.

We can use these results to derive some robust results among the individual technologies. Two of these technologies have robust results. Bio-electricity has a high investment in every non-dominated portfolio,

so Bio-Electricity-High is in the core. This technology appears to be good regardless of what probability distribution is used to evaluate it. Nuclear has either a Mid or High investment in every non-dominated portfolio, so Nuclear-Low is excluded. It is robust to invest in nuclear at least at the mid- level, regardless of the probability distribution used.

III.4.b Robustness Concepts

Figure 2 illustrates two ways to present the results, with different implications for robustness.

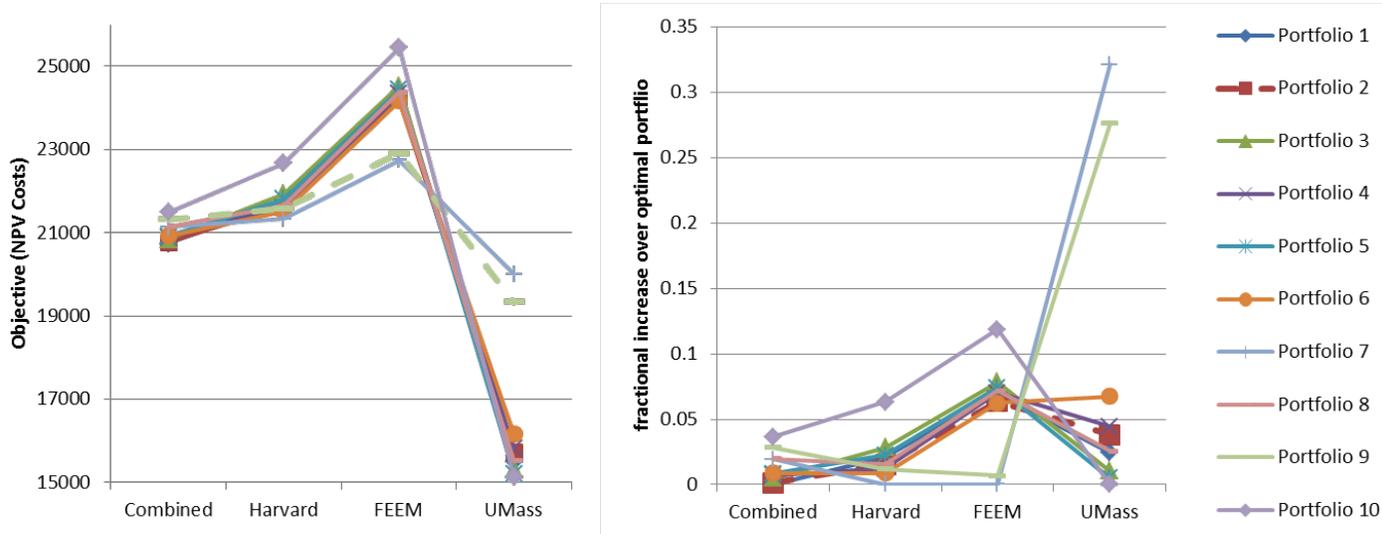


Figure 2: Left panel reports the absolute value of the NPV costs for each of the non-dominated portfolio (represented by different colors) calculated according to each of the four sets of beliefs. Right panel reports the relative increase in NPV costs with respect to the optimal portfolio for each of the non-dominated portfolios (represented by different colors) calculated according to each of the four sets of beliefs.

The left panel in Figure 2 reports absolute values of the objective, $H(x, \tau)$. The right panel in Figure 2 shows, for each team, the fractional increase above the individual team’s lowest cost. Mathematically, we show for each team τ :

$$\frac{H(x, \tau)}{\min_x H(x, \tau)} - 1$$

Both representations have benefits and drawbacks. The first highlights which beliefs lead to large absolute differences in objective value between the portfolios; and also allows decision makers to eliminate alternatives which do not satisfy an absolute threshold. This concept is most easily connected with the concept of the MiniMax. Portfolio 9, which has been dashed in this figure, would be the solution to the MiniMax. The least robust portfolio using that concept would be portfolio 10.

This representation has a drawback that is exacerbated when working with portfolios. This robustness concept, of avoiding high absolute costs, leads to an over-reliance on beliefs that are pessimistic. That is, the most robust portfolio will tend to be the optimal portfolio under the most pessimistic team. The problem is that there is no good reason to think that the most pessimistic team has any greater insight into the tradeoffs between portfolios and especially into the tradeoffs between individual alternatives. The representation on the right normalizes objective values so that each portfolio is compared to the portfolio that is best for that team, regardless of the absolute value of the objective. This concept is most easily connected with MiniMax Regret (indeed, Portfolio 2 -the dashed line in the right hand side panel of Figure 2- represents the solution to a MiniMax regret problem). The least robust portfolios using this concept are Portfolios 7, 9, and 10, since each of these has a high increase in cost, or regret, under at least one team’s beliefs. The drawback of this representation emerges when sets of beliefs lead to very low objective values in absolute terms. In this case, a small increase measured in absolute value would be magnified as a percentage increase, thus leading to an over-reliance on the most optimistic team. In this case, by looking at the first panel, it is clear that this is not the case: portfolios 7 and 9 are clearly significant outliers in an absolute sense for the UMass team.

Table 4 shows the portfolios that remain after we remove 7,9, and 10. We now explore this subset of portfolios for robust alternative across them, finding robust results for all technologies: (i) Solar-High is exterior, meaning that either a low or a medium investment is robustly non-dominated under all distributions; (ii) Nuclear-Low is exterior: either a mid or a high investment is robustly non-dominated; (iii) BioFuels-Low is exterior, as again either a mid or high investment is robustly non dominated; (iv) Bio-Electricity-High stays in the core; (v) CCS – mid is in the core, as CCS has a mid investment in every robustly non-dominated portfolio.

Portfolios	Technologies					Objectives ENPV (cost in billions of \$2005)			
	Solar	Nuc	BF	BE	CCS	Combined	Harvard	FEEM	U Mass
1	Low	High	High	High	Mid	20736	21770	24327	15509
2	Low	Mid	High	High	Mid	20768	21654	24188	15720
3	Low	High	Mid	High	Mid	20838	21929	24525	15301
4	Mid	High	High	High	Mid	20889	21588	24345	15813
5	Low	Mid	Mid	High	Mid	20912	21806	24434	15213
6	Mid	Mid	High	High	Mid	20922	21513	24163	16162
8	Mid	Mid	Mid	High	Mid	21144	21659	24379	15528

Table 4: Non-dominated portfolios with least robust portfolios (# 7, 9, 10) removed.

In fact, there are only eight possible portfolios that contain a combination of mid or low for solar, mid or high for nuclear and biofuels, high for Bio-electricity and mid for CCS. Of these eight portfolios, all but one are in the non-dominated group in Table 3.

These results imply that, given the GCAM model and a stabilization objective of 450 ppmeq by 2100, the policy recommendation would be to fund Bio-Electricity high, CCS mid, and some combination of solar, nuclear, and biofuels as recommended, depending on other concerns. For example, if budgets were tight, those three technologies could be funded at low, mid, mid. If nuclear is controversial, it could be funded at mid.

III.4.c Measures of Individual Alternatives

Finally, if a decision maker wanted a specific recommendation, one possibility is to identify the portfolio made up of projects that have Core Index greater than 0.5. Recall that the Core Index is the ratio of the number of non-dominated portfolios that contain a project with the total number of non-dominated portfolios. For example, Solar Low has a Core Index of $5/10=0.5$ among all non-dominated portfolios and $4/7=0.57$ among the robustly non-dominated portfolios. If we consider only the robustly non-dominated portfolios, then the portfolio with Core Indices greater than 0.5 for all of the technologies is portfolio 2, with investments in Solar Low, Nuclear Mid, Biofuels High, Bio-electricity High, and CCS Mid. If we consider the full set of Non-dominated portfolios, then the result remains the same, except for biofuels: no project has a CI of 0.5 or over and Mid and High are tied. This may indicate that biofuels is the particular technology which requires more investigation.

Although results are conditional to a specific assumption on climate policy and to the model used to represent technology implications for society, the key role for Bio-Electricity with CCS emerging from Table 4 has been widely documented as reported in (Clarke et al, 2014). Nevertheless, it will be critical to perform the full analysis using multiple integrated assessment models and multiple climate targets, to derive a more robust assessment of the future socio-economic value of technological improvements. As Table 2 shows, the R&D investment amounts vary considerably from technology to technology. For example, the “high” investment amounts for bio-electricity and biofuels are similar to the “mid” amounts for nuclear and CCS. Figure 3 compares the actual investments associated with the portfolio with the highest CI to those associated with the smallest and largest non-dominated portfolios. The

dollar investment level in all the technologies but solar is about even in the portfolio with the highest CI. The smallest non-dominated portfolio cuts investment in biofuels, while the largest non-dominated portfolio implies disproportionately larger investments in CCS and Nuclear.

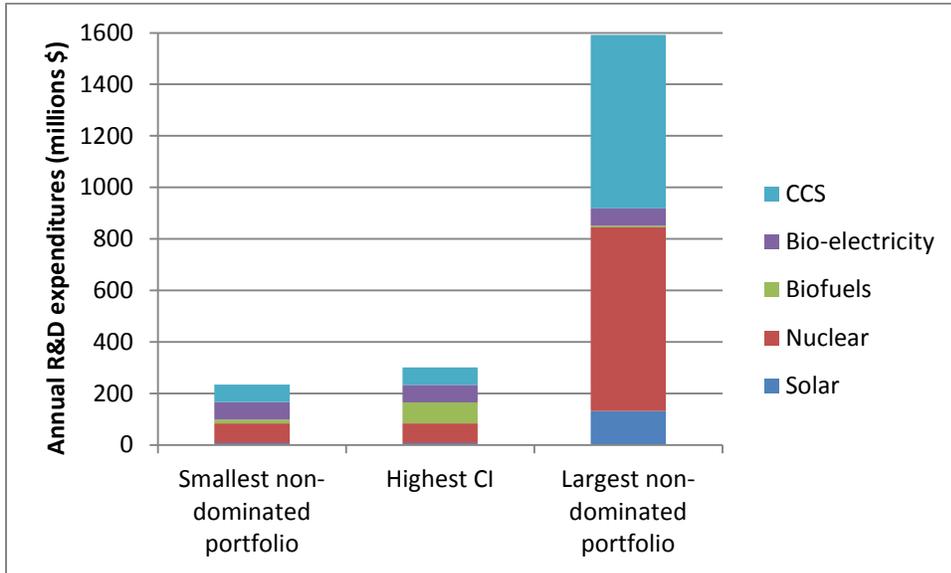


Figure 3: R&D Expenditure associate to various portfolios

IV. Flexibility of the Framework

Section III presented a proof of concept of a new method to aid decision processes in the face of deep uncertainty and conflicting beliefs. Here we illustrate the flexibility of this framework.

We note that there are many different types and sources of deep uncertainty in the climate change world. In this paper, we have specifically addressed multiple beliefs about one specific type of uncertainty: uncertainty over well-defined parameters (such as technology costs) represented by probability distributions. Another type of uncertainty is sometimes called “model uncertainty,” and refers to the uncertainty that is derived from the representation of processes in models. For example, in our analysis we have employed a single specific IAM, the GCAM model, to translate technology

parameters into societal costs and benefits. There exist a variety of IAMs that could be employed to provide the same analysis; this would likely result in different rankings over alternatives.

Our framework of belief dominance is flexible enough to allow for considering different models as another source of different beliefs. Reconsider our model presented in equation (5). Let m represent a particular model. The objective function becomes:

$$H(x; \tau, m) = \min_x \int \{TAC_m(\mathbf{z}, s) + \kappa B(\mathbf{x})\} f_\tau(\mathbf{z}; \mathbf{x}) d\mathbf{z}, \quad (8)$$

While we focused on the probability distribution f over the parameters \mathbf{z} , we are ultimately concerned with the implied probability distribution over the TAC. Individual values of the TAC depend on the outcomes of technological change, \mathbf{z} , and on the model used to estimate the TAC, m . Thus, the combination of a particular distribution over \mathbf{z} and a particular IAM produces a particular distribution over the TAC.

A portfolio x *belief dominates* x' if

$$H(\mathbf{x}; \tau, m) \geq H(\mathbf{x}'; \tau, m) \quad \forall \tau, m \quad (9)$$

Here a “belief” refers to a combination of τ and m that leads to a particular distribution over TAC.

Moreover, our framework is not limited to traditional portfolio problems, such as technology R&D. It can be applied more broadly to a wide range of applications, including a broader interpretation of climate change policy. Individual alternatives can include not only investments into energy technologies, but other technology policies, such as standards or subsidies, as well as other climate change policies, such as carbon taxes, carbon caps, international trade agreements, or near-term adaptation decisions. Uncertainties can include not only technological progress, but damage uncertainty, socio-economic uncertainties, and model uncertainty.

V. Conclusions

We present Robust Portfolio Decision Analysis as a promising approach to deal with problems of decision making in the face of deep uncertainty, i.e. situations characterized by significant disagreement about probability distributions over relevant outcomes. The two key aspects of our approach, building on the ideas of Robust Portfolio Modeling (RPM; Liesiö et al., 2007, 2008), are that (1) it allows us to

define non-dominated portfolios of strategies or decisions, in the face of multiple, conflicting beliefs over relevant outcomes; and (2) it allows us to derive insights and implications about individual strategies by looking at the portfolio-level results.

We demonstrate our approach on the specific case of designing a portfolio of publically-funded research and development investments in future energy technologies. From this, we find some common ground among the divergent expert beliefs, namely that two investments (mid-level in CCS and high-level in Bioelectricity) are in all of the most robust non-dominated portfolios, when the climate goal is stringent and GCAM defines the economic benefits. Policy negotiators could build on this common ground, incorporate non-quantifiable criteria, and perhaps commission more information where it is most likely to impact decisions, such as into biofuels.

This method presents innovative and useful elements that can generate important steps forward in the decision-making process of several societal problems that are affected by deep uncertainty. It does not ignore knowledge, nor does it ignore uncertainty and disagreement. It has promise to provide analytically rigorous support to decision making under deep uncertainty while preserving flexibility for decision makers. The combination of finding common ground and preserving flexibility may help to catalyze difficult dialog.

APPENDIX

Here we reprint a figure from Baker et al 2015, illustrating the standardized data set of four sets of beliefs over eight technology parameters.

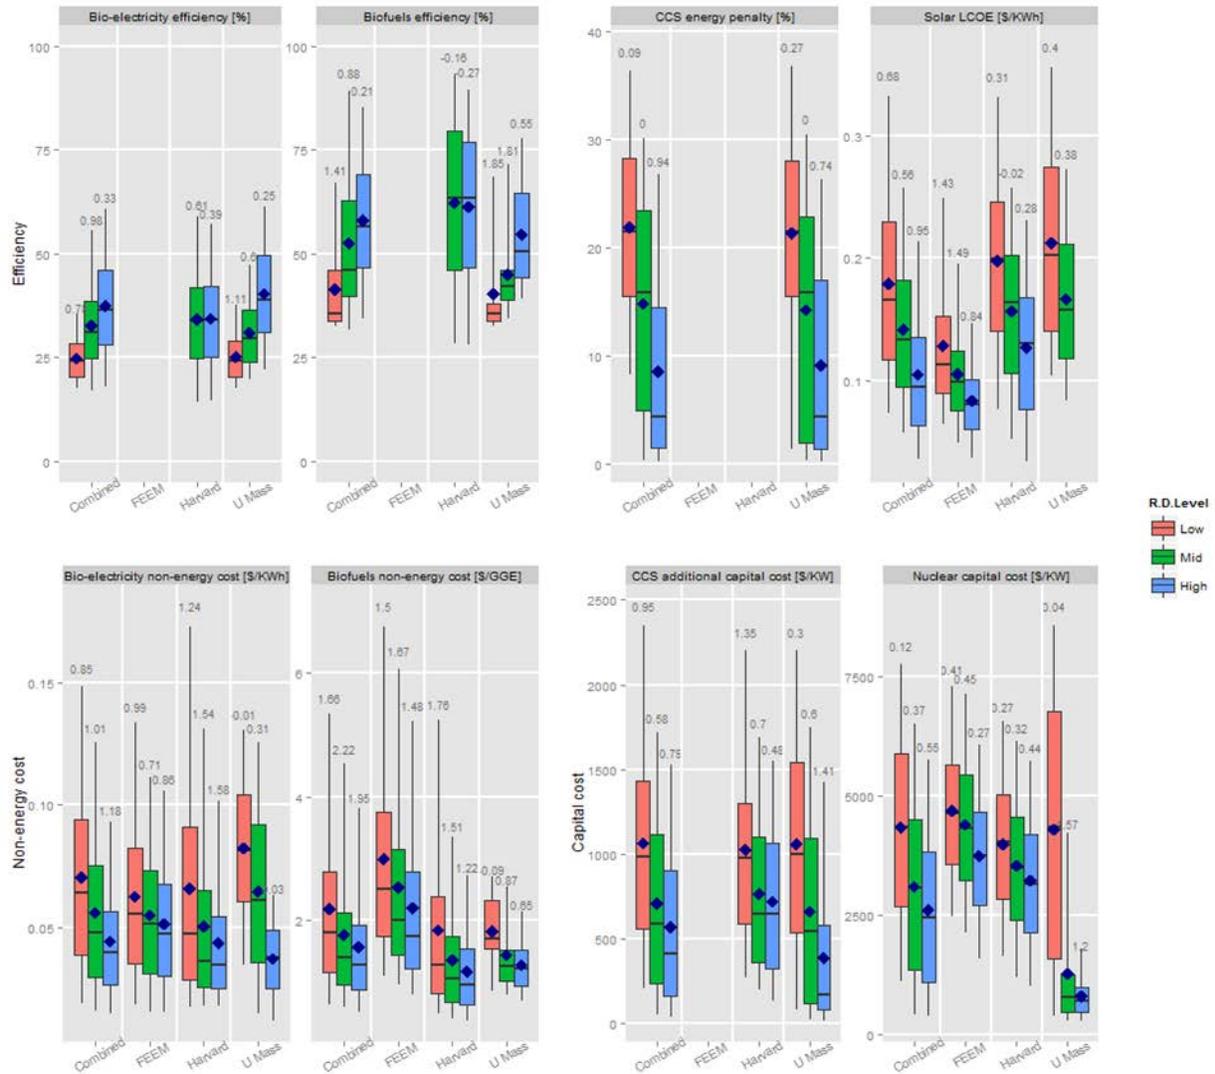


FIGURE A1: REPRINTED FROM [BAKER ET AL 2015] (NEED PERMISSION). : 2030 costs and efficiency elicitation results across studies and R&D levels. We show the combined distribution of the three studies using equal weights (“Combined”), the FEEM aggregate, the Harvard aggregate, and the UMass aggregate and technologies by R&D level (Low, Mid, and High). The box plots show the 5th, 25th, 50th, 75th, and 95th percentiles for each of the distributions, the diamond the mean value, and the black number the skewness of the distribution.

THEOREM. LET F_1, F_2, \dots, F_N BE A FINITE SET OF BELIEFS, REPRESENTED AS PROBABILITY DISTRIBUTIONS, AND LET Φ BE THE SET OF CONVEX COMBINATIONS OF THE BELIEFS SUCH THAT $\Phi = \{F = \lambda_1 F_1 + \lambda_2 F_2 + \dots + \lambda_N F_N \mid \text{FOR ALL } \lambda_i, i=1, \dots, N \text{ SUCH THAT } \lambda_i \geq 0, \sum_i \lambda_i = 1\}$. ALTERNATIVE X BELIEF DOMINATES X' FOR ALL DISCRETE BELIEFS $i=1, \dots, N$, IF AND ONLY IF X BELIEF DOMINATES X' OVER Φ .

PROOF OF THEOREM 1:

Assume that X belief dominates X' for all discrete beliefs i :

$$\int U(x; z) f_i(x; z) dz \geq \int U(x'; z) f_i(x'; z) dz$$

If we multiply each of the dominance inequalities by the associated term λ_i and sum, we get:

$$\sum_i \lambda_i \int U(x; z) f_i(x; z) dz \leq \sum_i \lambda_i \int U(x'; z) f_i(x'; z) dz$$

By linearity of the integral this implies:

$$\int [U(x; z) \sum_i \lambda_i f_i(x; z)] dz \leq \int [U(x'; z) \sum_i \lambda_i f_i(x'; z)] dz$$

The converse, that X belief dominates X' over the discrete set if it dominates it over the full set Φ , is trivial.

REFERENCES

- Al-Najjar, N. I. and J. Weinstein (2009). The ambiguity aversion literature: A critical assessment, *Economics and Philosophy* 25/03 (2009): 249-284.
- Anadon, L. et al., 2012. Expert judgments about RD&D and the future of nuclear energy. *Environ. Sci. Technol.*, Volume 46, pp. 11497-504.
- Anadón, L., Chan, G. & Lee, A., 2014a. Transforming U.S. Energy Innovation. In: L. Anadón, G. Chan & A. Lee, eds. *Expanding and Better Targeting U.S. Investment in Energy Innovation: an Analytical Approach*. Cambridge, U.K., and New York, NY, USA: Cambridge University Press.
- Athanassoglou, S. and V. Bosetti (2015). Setting environmental policy when experts disagree, *Environmental and Resource Economics* 61/4, 495-516.
- Baker, E., V. Bosetti, L. D. Anadon, M. Henrion, and L. A. Reis (2015). Future costs of key green energy technologies: Harmonization and aggregation of energy technology expert elicitation data. *Energy Policy* 80 C, 219–232.
- Baker, E., Chon, H. & Keisler, J., 2009a. Carbon capture and storage: combining economic analysis with expert elicitations to inform climate policy. *Climatic Change*, Volume 96, pp. 379-408.
- Baker, E., Chon, H. & Keisler, J., 2009b. Advanced solar R&D: Combining economic analysis with expert elicitations to inform climate policy. *Energy Economics*, Volume 31, pp. S37-S49.
- Baker, E., Chon, H. & Keisler, J. M., 2008. *Advanced Nuclear Power: Combining Economic Analysis with Expert Elicitations to Inform Climate Policy*. 08 August.
- Baker, E. & Keisler, J., 2011. Cellulosic biofuels Expert views on prospects for advancement. *Energy*, Volume 36, pp. 595-605.
- Baker, E. and E. Regnier (2015). Ambiguity aversion as a normative standard for policy-makers: The hazard of a framing effect, working paper.
- Baker, E. and S. Solak (2014). Management of energy technology for sustainability: How to fund energy technology R&D, *Production and Operations Management*, 23/3, 348-365.
- Ben-Haim, Y. (2004). Uncertainty, Probability and Information-Gaps. *Reliability Engineering & System Safety*, 85(1-3):249-266
- Ben-Tal, Aharon, Dimitris Bertsimas, and David B. Brown. "A soft robust model for optimization under ambiguity." *Operations research* 58.4-part-2 (2010): 1220-1234.
- Berger, L., J. Emmerling, and M. Tavoni (2015). Managing catastrophic climate risks under model uncertainty aversion. Accepted for publication in *Management Science*.
- Bertsimas, Dimitris, David B. Brown, and Constantine Caramanis. "Theory and applications of robust optimization." *SIAM review* 53.3 (2011): 464-501.
- Brown, C., Y. Ghile, M. Laverty, and K. Li (2012). Decision Scaling: Linking Bottom-up Vulnerability Analysis with Climate Projections in the Water Sector. *Water Resources Research*, 48, W09537.
- Bosetti, V., Catenacci, M., Fiorese, G. & Verdolini, E., 2012. The future prospect of PV and CSP solar technologies: An expert elicitation survey. *Energy Policy*, Volume 49, pp. 308-317.
- Box, G. E. P. and N.R. Draper (1987). *Empirical Model Building and Response Surfaces*, John Wiley & Sons, New York, NY.
- Caldeira, K., Jain, A. K., & Hoffert, M. I. (2003). Climate sensitivity uncertainty and the need for energy without CO2 emission. *Science*, 299(5615), 2052-2054.

- Catenacci, M. V. E., Bosetti, V. & Fiorese, G., 2013. Going electric: Expert survey on the future of battery technologies for electric vehicles. *Energy Policy*, Volume 61, p. 403–413.
- Center for Research on Environmental Decisions and ecoAmerica. (2014). *Connecting on Climate: A Guide to Effective Climate Change Communication*. New York and Washington, D.C.
- Chan, G., Anadón, L., Chan, M. & Lee, A., 2011. Expert elicitation of cost, performance, and RD&D budgets for coal power with CCS. *Energy Procedia*, p. 2685–2692.
- Chen, X., Sim, M., & Sun, P. (2007). A robust optimization perspective on stochastic programming. *Operations Research*, 55, 1058–1071.
- Clarke L., K. Jiang, K. Akimoto, M. Babiker, G. Blanford, K. Fisher-Vanden, J.-C. Hourcade, V. Krey, E. Kriegler, A. Löschel, D. McCollum, S. Paltsev, S. Rose, P.R. Shukla, M. Tavoni, B. van der Zwaan, and D. P. van Vuuren. 2014. "2014: Assessing Transformation Pathways." In *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge, United Kingdom and New York,: Cambridge University Press.
- Clemen, R. T., & Winkler, R. L. (1999). Combining probability distributions from experts in risk analysis. *Risk analysis*, 19(2), 187-203.
- Cooke, Roger M., and Louis LHJ Goossens. 2008. "TU Delft Expert Judgment Data Base." *Reliability Engineering & System Safety* 93 (5): 657–74.
- Drouet, L., V. Bosetti, M. Tavoni (2015), Selection of climate policies under the uncertainties outlined in IPCC AR5, *Nature Climate Change* 5, 937–940 (2015)
- Fiorese, G., Catenacci, M., Verdolini, E. & Bosetti, V., 2013. Advanced biofuels: Future perspectives from an expert elicitation survey. *Energy Policy*, Volume 56, p. 293–311.
- Gabrel, Virginie, Cécile Murat, and Aurélie Thiele. "Recent advances in robust optimization: An overview." *European Journal of Operational Research* 235.3 (2014): 471-483.
- Gilboa, I., A. Postlewaite, and D. Schmeidler (2009). Is it always rational to satisfy Savage's axioms? *Economics and Philosophy* 25/3, 285–296.
- Gilboa, I. and D. Schmeidler. 1989. Maxmin expected utility with nonunique prior. *Journal of Mathematics and Economics* 18(2): 141–53.
- Gregory, R., & Keeney, R. L.. (1994). Creating Policy Alternatives Using Stakeholder Values. *Management Science*, 40(8), 1035–1048. Retrieved from <http://www.jstor.org/stable/2633092>
- Hadka, D., J. Herman, P.M. Reed, K. Keller (2015) An open source framework for many-objective robust decision making *Environ. Model. Softw.* (in press).
- Heal, G. and, A. Millner (2014). Reflections: Uncertainty and decision making in climate change economics, *Review of Environmental Economics and Policy* 8/1, 120-137.
- Herman, J. D., P. M. Reed, H. B. Zeff, and G. W. Characklis (2015). How should robustness be defined for water systems planning under change? *Journal of Water Resources and Management*, 04015012, doi:10.1061/(ASCE)WR.1943-5452.0000509.
- Hora, Stephen C., Benjamin R. Fransen, Natasha Hawkins, and Irving Susel. 2013. "Median Aggregation of Distribution Functions." *Decision Analysis* 10 (4): 279–91.
- IPCC, 2014: *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* [Edenhofer, O., R. Pichs Madruga, Y. Sokona, E. Farahani, S. Kadner, K. Seyboth, A. Adler, I. Baum, S. Brunner, P.

- Eickemeier, B. Kriemann, J. Savolainen, S. Schlömer, C. von Stechow, T. Zwickel and J.C. Minx (eds.]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- Kalra, Nidhi; Hallegatte, Stephane; Lempert, Robert; Brown, Casey; Fozzard, Adrian; Gill, Stuart; Shah, Ankur. 2014. *Agreeing on robust decisions : new processes for decision making under deep uncertainty*. Policy Research working paper ; no. WPS 6906. Washington, DC: World Bank Group.
- Kasprzyk, J. R., S. Nataraj, P. M. Reed, and R. J. Lempert (2013). Many objective robust decision making 897 for complex environmental systems undergoing change. *Environmental Modelling & Software*, 42:55-71.
- Keeney, R. (1996). *Value-Focused Thinking: A Path to Creative Decisionmaking*. Harvard University Press, Cambridge, MA.
- Keller, K., B. M. Bolker, and D. F. Bradford (2004). Uncertain climate thresholds and optimal economic growth, *Journal of Environmental Economics and Management* 48/1, 723-741.
- Kelly, David L., and Charles D. Kolstad (1999). Bayesian learning, growth, and pollution, *Journal of Economic Dynamics and Control* 23.4 (1999) 491-518.
- Kim, S. H., Edmonds, J., Lurz, J., Smith, S. J., & Wise, M. (2006). The O bj ECTS framework for integrated assessment: Hybrid modeling of transportation. *The Energy Journal*, 63-91.
- Klibanoff, P., Marinacci, M., & Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6), 1849-1892.
- Lempert, R. J. (2002). A New Decision Sciences for Complex Systems. *Proceedings of the National 915 Academy of Sciences*, 99, 7309-7313.
- Lempert, R. and M. Collins (2007). Managing the risk of uncertain threshold responses: Comparison of robust, optimum, and precautionary approaches, *Risk Analysis* 27/4, 1009-1026.
- Lempert, R. J., Groves, D. G., Popper, S. W., & Bankes, S. C. (2006). A general, analytic method for generating robust strategies and narrative scenarios. *Management science*, 52(4), 514-528.
- Lichtendahl, Kenneth C. Jr, Yael Grushka-Cockayne, and Robert L. Winkler. 2013. "Is It Better to Average Probabilities or Quantiles?" *Management Science* 59 (7): 1594–1611.
- Liesiö, J., P. Mild, and A. Salo (2007). Preference Programming for Robust Portfolio Modeling and ProjectSelection, *European Journal of Operational Research* 181/3, 1488-1505.
- Liesiö, J., P. Mild, and A. Salo (2008). Robust Portfolio Modeling with Incomplete Cost Information and Project_Interdependencies, *European Journal of Operational Research* 190/3, 679-695.
- Liesiö, J. and A. Salo (2012). Scenario-based portfolio selection of investment projects with incomplete probability and utility information, *European Journal of Operational Research*, 217/1, 162-172.
- Mansbridge, J., Warren, M., Bächtiger, A., Cameron, M. A., Chambers, S., Ferejohn, J., ... & Tamir, Y. (2013). *Negotiating Agreement in Politics*. American Political Science Association.
- McInerney, D., R. Lempert, and K. Keller (2012). What are robust strategies in the face of uncertain climate threshold responses? *Climatic Change* 112/3-4, 547-568.
- Morgan, M. Granger, and M. Henrion. "Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis Cambridge University Press." New York, New York, USA, 1990.
- Millner, A., S. Dietz, and G. Heal (2013). Scientific ambiguity and climate policy, *Environmental and Resource Economics* 55/1, 21-46.
- Nordhaus, W. D. (2002). Modeling induced innovation in climate-change policy. *Technological change and the environment*, 9, 259-290.

- Oreskes, N., & Conway, E. M. (2011). *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming*. Bloomsbury Publishing USA.
- Popp, D. (2006). Innovation in climate policy models: Implementing lessons from the economics of R&D. *Energy Economics*, 28(5), 596-609.
- Ribas, G. P., Hamacher, S., & Street, A. (2010). Optimization under uncertainty of the integrated oil supply chain using stochastic and robust programming. *International Transactions in Operational Research*, 17, 777–796.
- Rosenhead, J., M. Elton, and S. K. Gupta (1972). Robustness and optimality as criteria for strategic decisions, *Operational Research Quarterly* 23/4, 413-431.
- Roser-Renouf, C., Maibach, E.W., Leiserowitz, A., & Zhao, X. (2014). The genesis of climate change activism: From key beliefs to political action. *Climatic Change*, 125(2), 163–178.
- Siniscalchi, M. (2009). Two Out Of Three Ain'T Bad: A Comment On "The Ambiguity Aversion Literature: A Critical Assessment". *Economics and Philosophy*, 25(03), 335-356.
- Tol, Richard S J (2013) Targets for global climate policy: an overview. *Journal of Economic Dynamics and Control*, 37 (5). pp. 911-928. ISSN 0165-1889
- Wallace, S. W. (2000). Decision making under uncertainty: Is sensitivity analysis of any use?. *Operations Research*, 48(1), 20-25.
- Webster, M., Jakobovits, L., & Norton, J. (2008). Learning about climate change and implications for near-term policy. *Climatic Change*, 89(1-2), 67-85.