Management of Energy Technology for Sustainability: How to Fund Energy Technology R&D

Erin Baker
Department of Mechanical and Industrial Engineering, College of Engineering, University of Massachusetts, Amherst, MA 01003, edbaker@ecs.umass.edu, Phone: +1-413-545-0670, Fax: +1-413-545-1027

Senay Solak*
Department of Operations and Information Management, Isenberg School of Management, University of Massachusetts, Amherst, MA 01003, solak@isenberg.umass.edu, Phone: +1-413-545-5681, Fax: +1-413-545-3858

Operations management methods have been applied profitably to a wide range of technology portfolio management problems, but have been slow to be adopted by governments and policy makers. We develop a framework that allows us to apply such techniques to a large and important public policy problem: energy technology R&D portfolio management under climate change. We apply a multi-model approach, implementing probabilistic data derived from expert elicitations into a novel stochastic programming version of a dynamic integrated assessment model. We note that while the unifying framework we present can be applied to a range of models and data sets, the specific results depend on the data and assumptions used, and therefore may not be generalizable. Nevertheless, the results are suggestive, and we find that the optimal technology portfolio for the set of projects considered is fairly robust to different specifications of climate uncertainty, to different policy environments, and to assumptions about the opportunity cost of investing. We also conclude that policy makers would do better to over-invest in R&D rather than under-invest. Finally, we show that R&D can play different roles in different types of policy environments, sometimes leading primarily to cost reduction, other times leading to better environmental outcomes.

Key words: energy technology, R&D portfolio, climate change, public policy, stochastic programming
History: Received: September 2012; accepted: January 2103 by Daniel Guide after two revisions

---

* Corresponding author

1. Introduction

Climate change is one of the biggest public policy problems currently facing the world. It is a very difficult problem for a number of reasons, including the long time frames, the global nature of the problem, and the deep uncertainty surrounding it. It is becoming clear that rapid technological change will be necessary in order to limit climate change in a way that is consistent with sustainable economic growth and current policies (Nordhaus 2011). One way of supporting such rapid technological change is through government-supported research and development (R&D) investment. While governments around the world have supported R&D for a very long time, there has been recent interest in applying a scientific basis to their resource allocation (National Research Council 2007).

In this paper, we develop a framework that uses empirical data for the assessment of possible R&D policy choices for sustainability. More specifically, we address the following important public policy questions: Given the uncertainty defined by currently available data in future technological success and climate change, what energy technology investment policies will maximize expected social welfare? And how do optimal investment policies differ under alternative strategies proposed for dealing with climate change, such as those suggested by Al Gore (Gore 2007), the Stern report (Stern 2007), and the Kyoto Protocol?

In order to address these questions and provide policy insights, we develop a multi-step multi-model approach involving a dynamic and stochastic R&D portfolio decision process. While doing so, we combine methods from multiple strands of research in operations management, including elicitation based decision analysis, stochastic programming, microeconomics, and computational economic analysis. In the remainder of this section we present the general framework for our problem, describe the research in the area, and discuss how our analysis and findings contribute to the existing literature.

1.1. General Decision Framework

There are two key near-term societal responses to climate change. The first and most direct is *abatement*, that is to reduce emissions of the greenhouse gases that are causing climate change to a level below what they would otherwise be. Examples of questions related to this response would be the determination of the optimal path of emissions in future years, emissions allocations, or the level of a carbon tax. A second response to climate change is to invest in *energy technology R&D* so that emissions abatement will be less costly in the future. A given emissions path influences the set of technologies society would like to have in the economy, and the set of technologies actually available influences the optimal level of emissions reductions. We explicitly recognize and model this interdependency as part of our analysis in this paper. Specifically, we simultaneously determine the optimal investment in a portfolio of technology R&D projects and the optimal emissions path so that the expected societal costs of climate change are minimized. Our analysis is a global one in that climate change is a global problem, with worldwide emissions affecting all parts of the globe. On the other hand, the R&D project data is based on U.S. government investment options.

The decision process we consider for our R&D investment optimization framework consists of two distinct but interlinked decision stages. These correspond to *near-term decisions* to be made over the next fifty years under climate change and technological uncertainty, and *long-term decisions* to be made after more information on uncertainties becomes available after the fifty year period. The near-term decisions are how much to invest in which technologies and the level of short term abatement. These decisions are made under the uncertainty of technical success which is dependent on the projects that are funded in each technology category, where success for a project means that a particular goal has been met. Hence, the uncertainty over technological success is endogenous: it depends on the decisions that will be made. Each technological success realization has a specific implication for future abatement costs, meaning that technological success will impact the costs of reducing emissions. The second stochastic characterization in our framework is an exogenous one and corresponds to the damages due to climate change. The uncertainty in these damages is represented through the parameters of a damage function, where the damage function depends on the stock of emissions in the atmosphere. The second stage decisions, which involve long-term abatement decisions, will be determined after information about future abatement costs and the damages becomes available. The objective for the overall decision problem is to maximize expected total social utility over the entire planning horizon involving the next four centuries. Surrounding all these decisions is the policy framework that defines the boundaries and limits for the decision making process.

1.2. Relevant Literature

Previous approaches to addressing climate change policy have included a great deal of theoretical work looking at how the optimal near-term policy changes with different characterizations of uncertainty (see Baker (2009) for a review). These studies, however, do not involve R&D decisions and use purely illustrative probability distributions to represent uncertainty. While Baker and Shittu (2008) review a set of papers that study R&D decisions in the face of climate and/or technology uncertainty, these papers are again based on illustrative distributions, and moreover, they consider only one technology at a time.

There are few papers that study the impact of uncertainty on a *portfolio* of energy technologies (see Baker and Solak (2011) for a review, including those that consider learning-by-doing rather than R&D). The one study we know of, Blanford (2009), uses illustrative probability distributions and simply assumes there are decreasing returns to scale in R&D investment. Our study differs in that we use empirical probability distributions obtained through expert elicitations within a comprehensive stochastic portfolio model we develop.

While Baker and Solak (2011) also describe a stochastic R&D optimization model based on elicited numerical data, it is a simplified model that represents the economy and the impacts of climate change with a single equation and two planning periods. Due to this simple structure it is unable to consider multiple policy frameworks. On the other hand, the insights from this model serve as input to the comprehensive analysis performed in this paper.

On the other end of the spectrum from the theoretical analysis is a body of work based on technologically detailed Integrated Assessment Models (IAMs)[1], which integrate economic models with climate models in order to provide policy relevant insights (Clarke et al. 2008, 2009). While these analyses provide important insights into the value of technology in society's response to climate change, they do not explicitly incorporate uncertainty or address the question of the optimal R&D policy. One exception is a recent analysis by Anadon et al. (2011), where the authors combine empirical data with an IAM-based analysis to perform portfolio optimization. Unlike our study, however, their data do not differentiate between different projects within technology categories, and the optimization itself is not stochastic – they consider only the most likely outcome for any given R&D investment. Yet, a recent study by the National Academy suggests that uncertainty be explicitly included in the U.S. Department of Energy decisions about investments into R&D (National Research Council 2007). Thus, our study is unique in explicitly incorporating uncertainty and stochastic R&D optimization within a detailed IAM-based analysis, while maintaining tractability in the resulting model.

The remainder of this paper is structured as follows. In Section 2 we describe the general structure for our modeling approach. The components of the model are developed in Section 3, and stochastic optimization procedures are described in Section 4. The experimental setup for policy analysis is discussed in Section 5, while the numerical results and their implications are described in Section 6. Finally, in Section 7 we provide a summary of our conclusions.

[1] Acronyms used throughout the paper are summarized in Appendix S1 for reference purposes.

2. Integrated R&D and Abatement Policy Optimization Model

The stochastic optimization problem representing the decision process described in Section 1.1 involves the determination of an optimal portfolio of technology investments and an abatement policy such that the expected total social utility over the planning horizon is maximized. For a general mathematical representation of this problem, we first let the investment decisions be denoted by a vector $\boldsymbol{\Upsilon}$, and the near-term and long-term emissions abatement decisions by vectors $\boldsymbol{\mu}_N$ and $\boldsymbol{\mu}_L$, respectively[2]. Note that abatement is defined as the fraction of emissions reduced below business-as-usual level. Similar to the notation used for abatement, the vectors $\boldsymbol{y}_N$ and $\boldsymbol{y}_L$ represent the total output of goods and services in near-term and long-term, while $\boldsymbol{\tau}_N$ and $\boldsymbol{\tau}_L$ are the atmospheric temperatures. In addition, the vectors $\boldsymbol{x}_N$ and $\boldsymbol{x}_L$ are used to denote the set of other decision variables in each stage that define the relationships between climate change, economy and social utility. We represent the two critical functions in our framework, namely the abatement cost function and the damage function, as $c(\boldsymbol{\mu}_s)$ and $D(\boldsymbol{\tau}_s)$, $s = N, L$, indicating their dependence on abatement decisions and atmospheric temperatures, respectively. Moreover, we represent uncertainty about technology and climate change information through the set Ω of all possible scenarios, with $\omega \in \Omega$ representing a single scenario. The probability of occurrence for each scenario is denoted by p^ω. Given this representation, the integrated R&D and abatement policy optimization model is stated in general form as follows:

$$\max_{\boldsymbol{\Upsilon},\boldsymbol{\mu}_N,\boldsymbol{y}_N,\boldsymbol{\tau}_N,\boldsymbol{x}_N,\boldsymbol{\mu}_L^\Omega,\boldsymbol{y}_L^\Omega,\boldsymbol{\tau}_L^\Omega,\boldsymbol{x}_L^\Omega} U_N(\boldsymbol{\Upsilon},\boldsymbol{\mu}_N,\boldsymbol{y}_N,\boldsymbol{\tau}_N,\boldsymbol{x}_N) + \sum_{\omega\in\Omega} p^\omega U_L(\boldsymbol{\mu}_L^\omega,\boldsymbol{y}_L^\omega,\boldsymbol{\tau}_L^\omega,\boldsymbol{x}_L^\omega,\omega) \tag{1}$$

$$\text{s.t. } \boldsymbol{y}_N = g(c(\boldsymbol{\mu}_N), D(\boldsymbol{\tau}_N), \boldsymbol{x}_N) \tag{2}$$

$$\boldsymbol{y}_L^\omega = g(c(\boldsymbol{\mu}_L^\omega), D(\boldsymbol{\tau}_L^\omega), \boldsymbol{x}_N, \boldsymbol{x}_L^\omega, \omega) \qquad \forall\omega \tag{3}$$

$$\boldsymbol{G}_N(\boldsymbol{\Upsilon},\boldsymbol{\mu}_N,\boldsymbol{y}_N,\boldsymbol{\tau}_N,\boldsymbol{x}_N) \le \boldsymbol{b}_N \tag{4}$$

[2] The notation used throughout the paper is summarized in Appendix S2 for reference purposes.

$$G_L(\boldsymbol{\Upsilon}, \boldsymbol{\mu}_N, \boldsymbol{y}_N, \boldsymbol{\tau}_N, \boldsymbol{x}_N, \boldsymbol{\mu}_L^\omega, , \boldsymbol{y}_L^\omega, \boldsymbol{\tau}_L^\omega, \boldsymbol{x}_L^\omega, \omega) \leq \boldsymbol{b}_L^\omega \qquad \forall \omega \quad (5)$$

where the superscript in $\boldsymbol{\mu}_L^\omega, \boldsymbol{y}_L^\omega, \boldsymbol{\tau}_L^\omega$, and $\boldsymbol{x}_L^\omega$ denotes the dependence of these decision variables on the realized values of the stochastic parameters for each scenario, while the superscript in $\boldsymbol{\mu}_L^\Omega, \boldsymbol{y}_L^\Omega, \boldsymbol{\tau}_L^\Omega$, and $\boldsymbol{x}_L^\Omega$ implies that the optimization is performed over the corresponding variables in all scenarios, as we let $\boldsymbol{\mu}_L^\Omega = \langle \boldsymbol{\mu}_L^1, \ldots, \boldsymbol{\mu}_L^{|\Omega|} \rangle$, $\boldsymbol{y}_L^\Omega = \langle \boldsymbol{y}_L^1, \ldots, \boldsymbol{y}_L^{|\Omega|} \rangle$, $\boldsymbol{\tau}_L^\Omega = \langle \boldsymbol{\tau}_L^1, \ldots, \boldsymbol{\tau}_L^{|\Omega|} \rangle$, and $\boldsymbol{x}_L^\Omega = \langle \boldsymbol{x}_L^1, \ldots, \boldsymbol{x}_L^{|\Omega|} \rangle$. The function $g(\cdot)$ in constraints (2) and (3) represents the relationship between output, climate changes damages and abatement of emissions.

The functions $U_s(\cdot)$, $s = N, L$ represent total social utility over the near and long-term stages, each of which consists of multiple time periods. Hence, objective (1) involves maximization of the summation of the near-term and the expected long-term social utilities, where the expectation is defined over all possible scenarios. Equations (2) and (3) represent the relationship between output of goods and services, climate change damages and abatement costs for the near-term and long-term decision problems, respectively. Similarly, the constraint sets (4) and (5) correspond to the relationships defining the interplay between climate change, economy and social utility. Note that the second stage constraints (3) and (5), and thus the optimal long-term abatement policy, is dependent on the technology investments $\boldsymbol{\Upsilon}$, near-term abatement $\boldsymbol{\mu}_N$, the output levels $\boldsymbol{y}_N$, the atmospheric temperatures $\boldsymbol{\tau}_N$, and other related decisions $\boldsymbol{x}_N$ made in the initial decision stage.

There are several challenges, however, that need to be overcome to fully implement this model as a valid energy technology R&D policy analysis. These involve the development of functional representations and inputs for this general problem structure, as well as methodological integration and implementation within a tractable stochastic optimization framework. We address these challenges by seeking answers to the following questions through a multi-step multi-model process: (1) Step 1: Modeling of the investment options for decision vector $\boldsymbol{\Upsilon}$: What technology projects should be considered? What are the return characteristics for these technology projects?; (2) Step 2: Modeling of the social utility functions $U_s(\cdot)$, constraint sets $\boldsymbol{G}_s(\cdot)$, $s = N, L$, and the output function $g(\cdot)$: How should the interplay between social utility, climate change and the economy be modeled? How does R&D investment impact these relationships?; (3) Step 3: Modeling of uncertainty for scenario set Ω: How should the uncertainty in climate change damages and the uncertainty in R&D-induced technical change be modeled?; (4) Step 4: Implementation and solution of the stochastic optimization problem (1)-(5): How should different modeling components be integrated and implemented under a tractable stochastic optimization framework?

In the next two sections we describe how we address these issues to identify policy results for energy technology R&D under climate change. The first two steps listed above correspond to general *model component development*, and are discussed in Section 3. The last two steps, which involve *uncertainty modeling and stochastic optimization*, are described in Section 4.

3. Model Component Development

We discuss model component development separately for each of the first two steps listed in Section 2 above. For some parts of these steps we utilize results from relevant modeling and analysis efforts that are described in detail in other studies. In the descriptions below, we provide references to these studies and also discuss how we relate the results from these models to our integrated R&D and abatement policy optimization framework.

3.1. Step 1: Modeling of the Investment Options

Identification of technologies and projects. Our analysis considers investment options in three key technology areas: carbon capture and storage (CCS), nuclear fission, and solar photovoltaics.

While this does not cover the full portfolio of energy technologies, or even electricity technologies, it provides a good representation of the problem. Lewis and Nocera (2006) have pointed out in their analysis that these three technologies are the only ones with sufficient resources to provide the carbon-neutral energy needed to address the climate change problem. Thus, our work can provide specific insights into how to balance among three key technologies, as well as a framework that can be expanded into more technologies in the future as necessary. Each of the three key technology categories contains multiple research areas, or 'projects', in which R&D investments can be made. The research areas considered for each technology are listed in Appendix S3. These projects were chosen jointly with experts, with the aim of considering the projects in each technology that have the possibility of resulting in a breakthrough. Hence, they represent the most relevant investment options in each technology from a policy perspective.

Characterization of success probabilities for individual technology projects. The probability of success in each project is defined through expert elicitations. Expert elicitation is a formal process for quantifying an expert's judgement about uncertain quantities, and capturing those judgments in terms of probabilities that can be used in further analyses (Hora 2004). While expert elicitations are subject to a number of known biases (Tversky and Kahneman 1974), no other method exists that can be used to gain information about potential future breakthroughs in technologies. In fact, in a review of the climate change assessment of the Intergovernmental Panel on Climate Change, InterAcademy Council (2010) specifically suggests that " to inform policy decisions properly, it is important for uncertainties to be characterized and communicated clearly and coherently...[w]here practical, formal expert elicitation procedures should be used to obtain subjective probabilities for key results."

We base our model on the expert elicitations summarized in Appendix S3, and described in detail by Baker et al. (2008), Baker et al. (2009a) and Baker et al. (2009b). The elicitations assume that each project can be invested in at one of multiple potential levels, where investments include only U.S. government funding are measured based on net present value. Each project is also associated with specific endpoints or targets to be assessed, such as a given cost and efficiency level, which define 'success' for that project. The specific probabilities of success defined through the elicitations for different investment levels of each project reflect an aggregation of the individual experts' judgments. Expert input was also used to define both the funding levels and the endpoints to be assessed.[3]

3.2. Step 2: Modeling of the Social Utility Functions and the Constraint Sets

The modeling of the social utility functions $U_s(\cdot)$, the constraint sets $\boldsymbol{G}_s(\cdot)$ for $s = N, L$, and the output function $g(\cdot)$ involves multiple phases. First, the complex relationships between climate change, emissions abatement, economy and social utility are represented in basic form through a well-known IAM, namely the Dynamic Integrated Model of Climate and the Economy (DICE) (Nordhaus 1993). Then, these relationships are expanded to include R&D investments and the impact of resulting technical change. This requires the quantification of the impact of R&D on abatement cost functions, and the integration of this quantitative measure into the modeling framework as part of the constraint sets.

3.2.1. Basic Representation through the DICE Model DICE is a deterministic global optimal growth model that includes interactions between economic activities and the climate. The model covers a long planning horizon, typically around 400-600 years, in ten-year periods. In each period, economic output (measured as the gross domestic product (GDP)) is divided between consumption and investment in new capital, consistent with the standard optimal growth framework in economic analysis. DICE adds to this framework by modeling the emissions of greenhouse gases into the atmosphere as part of the production process. The general

[3] We note here that while other elicitation data exists on the three technologies, they are not applicable to the type of R&D portfolio analysis studied in this paper. See Appendix S4 for more details.

Parameters		Variables	
R_t :	utility discount factor for period t	o_t :	consumption of goods/services in period t
A_t :	level of total factor productivity in period t	y_t :	net output of goods/services in period t
S_t :	ratio of uncontrolled emissions to output in period t	y_t^g :	unadjusted output in period t
E_t :	emissions from deforestation in period t	k_t :	capital stock in period t
L_t :	population and labor input in period t	l_t :	investment in period t
β :	elasticity of marginal utility of consumption	e_t :	total carbon emissions in period t
γ :	elasticity of output with respect to capital	μ_t :	emissions abatement in period t
σ :	rate of depreciation of capital	τ_t :	atmospheric temperature in period t
	Functions	u_t :	social utility in period t
$H(\cdot)$:	function linking emissions to atmospheric temperature		
$c_D(\cdot)$:	abatement cost function in the DICE model		
$D_D(\cdot)$:	climate change damage function in the DICE model		

Table 1 List of parameters, variables, and functions used in the summary formulation of the DICE model

formulation of the DICE model is given in Nordhaus (2008), and is also summarized through equations (6)-(13) below. We will later extend constraints (8) and (11) in this formulation through inclusion of uncertainty and investment decisions. The summary formulation utilizes the parameters and variables shown in Table 1.

$$\max_{\boldsymbol{\mu},\boldsymbol{y},\boldsymbol{\tau},\boldsymbol{x}} \quad \sum_t R_t u_t \tag{6}$$

$$\text{s.t.} \quad u_t = L_t \frac{(\frac{o_t}{L_t})^{1-\beta} - 1}{1-\beta} \qquad \forall t \tag{7}$$

$$y_t = o_t + l_t \qquad \forall t \tag{8}$$

$$k_t = l_{t-1} + (1-\sigma)k_{t-1} \qquad \forall t \tag{9}$$

$$\tau_t = H\left(\tau_{t-1}, e_t\right) \qquad \forall t \tag{10}$$

$$y_t = \frac{1 - c_D\left(\mu_t\right)}{D_D\left(\tau_t\right)} y_t^g \qquad \forall t \tag{11}$$

$$y_t^g = A_t L_t^{1-\gamma} k_t^{\gamma} \qquad \forall t \tag{12}$$

$$e_t = S_t(1-\mu_t)y_t^g + E_t \qquad \forall t \tag{13}$$

Similar to the notation used in the general description in Section 2, the vector $\boldsymbol{x} = \langle \boldsymbol{o}, \boldsymbol{y}^g, \boldsymbol{k}, \boldsymbol{l}, \boldsymbol{e}, \boldsymbol{u} \rangle$ in objective (6) is used to refer to all other variables in the formulation, except for the abatement vector $\boldsymbol{\mu}$, output vector $\boldsymbol{y}$, and temperature vector $\boldsymbol{\tau}$. Key independent decision variables in the formulation are $\boldsymbol{\mu}$ and $\boldsymbol{o}$, while the others are dependent variables determined by the values of $\boldsymbol{\mu}$ and $\boldsymbol{o}$. Note that elements of $\boldsymbol{\mu}$, $\boldsymbol{y}$, $\boldsymbol{\tau}$ and $\boldsymbol{x}$ consist of both near-term and long-term decision variables as they include all periods, so no such distinction is made in the notation used. Moreover, the formulation does not involve R&D investment, which we later include through extensions of the given relationships.

In the formulation above, objective (6) ensures that policies are chosen to maximize the discounted sum of social utility u_t over time. Utility is based on per capita consumption o_t, as defined through the relationship in (7). This constraint defines utility in each period as an isoelastic function of consumption, where β is the calibrated elasticity parameter. Consumption is defined by equation (8) as the difference between output of goods/services y_t and the capital investment l_t. Capital balance relationship is represented through constraint (9). Constraint (10) represents a set of constraints that link greenhouse gas emissions e_t to the global temperatures τ_t. Note that the accumulation of greenhouse gases, especially carbon dioxide, affects welfare by increasing global temperatures. Thus, the representative function $H\left(\tau_{t-1}, e_t\right)$ is increasing in e_t. A key equation in the model is (11), which represents the relationship between output of goods/services y_t and the impacts of climate change. This representation involves the unadjusted output y_t^g in each period, which is determined by inputs of labor L_t and capital k_t as defined by (12). Note that the climate change damage function in DICE, i.e. $D_D(\tau_t)$ in the denominator of the right hand side of equation (11), is an increasing function of τ_t. This implies that a higher temperature negatively impacts the output y_t. In order to mitigate this effect, an abatement level μ_t can be chosen each period, which reduces emissions e_t below what would otherwise occur for a given production level. This relationship between abatement μ_t and emissions e_t is represented through constraint (13). While abatement has benefits, it is costly as the abatement cost function $c_D(\mu_t)$ in the numerator of (11) is increasing in μ_t. Hence, higher abatement reduces the output available for consumption or investment in every period. In other words, lower abatement positively impacts the output y_t. This tradeoff needs to be managed by choosing the best abatement effort μ_t in each period, as the optimal abatement path reflects a balance between benefits and costs.

Here we highlight two equations which we will use to incorporate R&D investments and stochasticity into our modeling framework. The first is equation (8), which shows how economic output y_t in each period is used. We will include R&D investments into the model through the modification of this constraint, which we describe in Section 3.2.3. The second equation we highlight is (11), which shows how the cost of abatement and the damages from climate change impact economic output y_t in each period t. The abatement cost function $c_D(\mu_t)$, for which a detailed description is provided in Appendix S5, will be revised to include the impact of technical change due to R&D investments. This procedure is also described in Section 3.2.3. On the other hand, $D_D(\tau_t) = 1 + \pi\tau_t^2$ in equation (11) represents the damages from climate change resulting from atmospheric temperature τ_t, where π is the damage parameter. We will take π to be a stochastic parameter in our model, representing the uncertainty in climate change damages, which we describe further in Section 4.1.1.

The DICE model formulation links to our general model (1)-(5) as follows: the utility function in (1) is represented by objective (6) and equation (7); equations (2)-(3) correspond to (11); and the general constraints (4)-(5) involve equations (8)-(10) and (12)-(13).

3.2.2. Quantifying the Impact of R&D on Abatement Costs The basic relationships represented through the DICE model need to be expanded to include R&D investment decisions and their impacts on other problem components. To achieve this, we first need to define measures that model how technical change resulting from R&D impacts abatement costs. We will then develop a procedure to integrate these quantified impacts into the modeling framework.

Recall that the second stage decision of ‘long-term abatement’, involves choosing an emissions abatement level for each period such that expected total social utility is maximized. Social utility is related to the cost of abatement, as higher costs would imply lower net economic output. On the other hand, the cost of abatement is dependent on the realized technological success in the invested R&D projects. Hence, we need to derive abatement cost functions under different technological success outcomes to implement in the model. We achieve this through the two phase process described below, which is also discussed by Baker and Solak (2011):

Deriving marginal abatement cost curves. Rather than deriving abatement cost functions directly for each technological success outcome, we first derive Marginal Abatement Cost Curves (MACs), which reflect the cost of reducing emissions by an additional tonne. We then use these curves to parameterize the impact of technical change on abatement cost functions.

We utilize the definitions of success described in Section 3.1 to derive MACs using a technologically detailed integrated assessment model, namely the Global Change Assessment Model (GCAM) (Brenkert et al. 2003, Edmonds et al. 2004). We use MACs rather than abatement cost curves for two reasons. First, the MAC is a key unit of analysis in our decision problem: as long as there is an interior solution to the second stage decision for a given realization of uncertain parameters, abatement will be optimized where the marginal cost of abatement is just equal to the marginal damages avoided by abatement. Second, the MAC is easy to generate from IAMs and is more likely to be consistent across IAMs than the abatement cost curve.

To derive the MACs, curves were generated in GCAM that relate levels of emissions reduction to carbon prices, thus approximating the marginal cost of abatement. This was done for each possible combination of projects assuming success in each project. A selection of the

MACs that were derived using GCAM are shown in Appendix S6, where the baseline MAC used to measure the impact of technical change refers to the case with no R&D-induced technical change as defined by Clarke et al. (2008). We demonstrate in Appendix S6 the impact of three projects (one from each technology, i.e. CCS, nuclear, and solar), if each of them were successful independently; as well as the MAC that is generated by GCAM if all three of the projects were successful simultaneously. The changes with respect to the baseline MAC are a combination of pivoting the curve clockwise around zero and shifting the entire curve down. CCS has more of a pivot, as it results in a nearly proportional reduction in the MAC. Nuclear, on the other hand, has more of a constant downward shift, with a much smaller pivot effect. Solar has a very similar qualitative impact as nuclear, although its overall impact is small.

Parameterizing the impact of R&D-induced technical change on abatement cost functions. The process described above resulted in numerical MAC curves for each combination of successful technology projects. In order to make this tractable and portable to our framework, the next step was to estimate parameters that quantify how each technology combination impacts the baseline MAC. To this end, for each combination of successful projects, we used equation (14) below to estimate a pivot vector $\boldsymbol{\alpha} = \langle \alpha_{CCS}, \alpha_{nuclear}, \alpha_{solar} \rangle$, with each component representing the pivoting effect of the corresponding technology based on the successful projects in that technology, as well as a shift parameter $h(\boldsymbol{\alpha})$. The shift parameter was represented as a function of $\boldsymbol{\alpha}$, because a unique shift effect was observed for every combination of values of α_i, $i = CCS, nuclear, solar$:

$$\widetilde{MAC}(\mu, \boldsymbol{\alpha}) = \prod_i (1 - \alpha_i) \, MAC(\mu) - h(\boldsymbol{\alpha}) \, MAC(0.5) \tag{14}$$

where $MAC(\mu)$ is the numerical baseline MAC from GCAM, with $\mu \in [0,1]$ denoting the level of abatement, and $\widetilde{MAC}(\mu, \boldsymbol{\alpha})$ is the estimated MAC after technical change, where technical change is represented through the vector $\boldsymbol{\alpha}$. The product term $\prod_i (1-\alpha_i)$ models the pivoting of the function $MAC(\mu)$ due to technical success, while $h(\boldsymbol{\alpha})$ models the corresponding shift. Notice that the shift $h(\boldsymbol{\alpha})$ is anchored on 50% abatement to make this representation portable from GCAM, in which the parameters were derived, to other modeling frameworks.

Let $\mathcal{S} = \cup_i \mathcal{S}_i$ refer to some given combination of successful technology projects, where $\mathcal{S}_i$ denotes the set of successful projects in technology i, $i = CCS, nuclear, solar$. The process for deriving the values of α_i and $h(\boldsymbol{\alpha})$ for any given set $\mathcal{S}$ was as follows. First, a project pivot parameter, denoted by α_{ij} was estimated using the generated MACs for each individual project $j \in \mathcal{S}_i$ in technology i. The values of these parameters, which vary depending on the level of success in a project, are listed in Appendix S7. Second, we make the assumption that, within any technology i, only the best project (the one with the greatest impact on the MAC) will impact the economy. Therefore, we define α_i as $\alpha_i = \max_j \{\alpha_{ij} : j \in \mathcal{S}_i\}$. Finally, for every combination of possible technological outcomes represented by the three α_i's for the three technologies, a shift parameter $h(\boldsymbol{\alpha})$ was estimated numerically.

The relationship in (14), which is defined over the MACs, needs to be transformed into a relationship over abatement costs for implementation in our model. We do this by integration and define a functional $\Phi(c(\mu), \boldsymbol{\alpha})$ that translates any generic baseline cost function (without technical change), denoted by $c(\mu)$, to a new abatement cost function with technical change. Specifically we define a functional $\Phi : \mathcal{F} \times \mathbb{R}^n \to \mathcal{F}$, where $\mathcal{F}$ is a set of functions and $\mathbb{R}^n$ is a set of real vectors, as follows:

$$\Phi(c(\mu), \boldsymbol{\alpha}) = \prod_i (1-\alpha_i)\, c(\mu) - h(\boldsymbol{\alpha})\, c(0.5)\, \mu \tag{15}$$

This can be used to model the impact of technical change on any abatement cost function.

3.2.3. Integrating R&D Investment and R&D-induced Technical Change into the Model The investment decisions for each technology and the parametric representation of the resulting impacts need to be integrated into the modeling framework through the introduction of new variables and constraints. We discuss this procedure here.

We noted in Section 3.2.2 that nuclear and solar have similar impacts, in that they both have strong shifts as well as pivots, as opposed to CCS which mainly has pivoting effects. Given this structure, we combine nuclear and solar into one category, calculating the resulting pivot values as $\alpha_2 = 1 - (1 - \alpha_{nuclear})(1 - \alpha_{solar})$. Thus, we represent R&D investments by using only two technology categories in the model ($i = 1$ for CCS and $i = 2$ for solar/nuclear).

Integrating R&D investment decisions into the model: Modification of constraint (8). We assume that R&D investment will take place over the next 50 years at a constant rate. In the basic representation in Section 3.2.1, output in each period is divided between consumption and investment in traditional capital as noted in equation (8). We extend this relationship to include R&D investment for periods $t \leq 5$ as follows. Note that each planning period corresponds to 10 years, i.e. $t = 5$ implies 50 years:

$$y_t = o_t + l_t + \kappa (\Upsilon_1 + \Upsilon_2)/5 \qquad \forall t \leq 5 \tag{16}$$

where κ is an opportunity cost multiplier, and Υ_i is the investment in each technology area i, with the index i corresponding to CCS and solar-nuclear as defined above. Hence, the total output in each period is either consumed, as represented by the variable o_t, or invested in traditional capital, as represented by the variable l_t, or invested in R&D, as represented by $\kappa (\Upsilon_1 + \Upsilon_2)/5$. While Υ_i is a decision variable, the opportunity cost κ is a parameter for which we perform sensitivity analysis as part of the numerical experiments described in Section 5.1.

Integrating R&D-induced technical change into the model: Modification of constraint (11). Following the cost function structure described by equation (15) in Section 3.2.2, we model the impact of technical change by altering constraint (11) as follows:

$$y_t = \frac{1 - \Phi\left(c_D\left(\mu_t\right), \boldsymbol{\alpha}\right)}{D_D\left(\tau_t\right)} y_t^g = \frac{1 - \prod_i (1-\alpha_i)\left[c_D\left(\mu_t\right) - h\left(\boldsymbol{\alpha}\right) c_D\left(0.5\right)\mu\right]}{D_D\left(\tau_t\right)} y_t^g \quad \forall t \qquad (17)$$

The model of technical change in equation (17), however, results in a non-convex model, involving multilinear-terms due to multiplication of α_i, which depend on the investment decision variables Υ_i. To deal with this, we estimate tight linear approximations to the actual functions. More specifically, we use our data on the α_i and $h\left(\boldsymbol{\alpha}\right)$ to estimate the two quantities of $(1-\alpha_1)(1-\alpha_2) \approx 1 - 0.8\alpha_1 - 0.92\alpha_2$ and $(1-\alpha_1)(1-\alpha_2)\, h\left(\boldsymbol{\alpha}\right) \approx 0.02 - 0.06\alpha_1 + 0.14\alpha_2$. Thus, we express the revised production function in our model for $t > 5$ as follows.

$$y_t = \frac{\left[1 - \left((1 - 0.8\alpha_1 - 0.92\alpha_2)c_D\left(\mu_t\right) - (0.02 - 0.06\alpha_1 + 0.14\alpha_2)c_D\left(0.5\right)\mu_t\right)\right]}{D_D\left(\tau_t\right)} y_t^g \qquad (18)$$

We show in Appendix S8 that modeling the pivot and shift effects through these linear approximations ensure that convexity of the optimization model is maintained in the extended stochastic model for the given practical bounds for variables.

4. Uncertainty Modeling and Stochastic Optimization

In this section, we describe how we model the stochastic structure in the energy technology R&D portfolio problem through a scenario set, and how we solve the resulting optimization model using stochastic programming.

4.1. Step 3: Modeling of Uncertainty

As noted in Section 1.1, our model involves two types of uncertainty: the exogenous uncertainty in climate change damages, and the endogenous uncertainty in technical change which is a function of the R&D investment decisions.

4.1.1. Modeling the Uncertainty in Climate Change Damages We model the uncertainty in climate change damages through a probabilistic characterization of the parameter π in the damage equation $D_D(\tau_t) = 1 + \pi {\tau_t}^2$. This is done by defining a three-point discrete probability distribution for π based on previous elicitations (Nordhaus 1994). As part of our analysis, we consider several risk cases for climate change, which are represented by different distributions of the parameter π. These risk cases and distributions are discussed in Section 5.3.

4.1.2. Modeling the Uncertainty in Technical Change In Section 3.1 we identified individual success probabilities for technology projects, and in Section 3.2.2 we modeled how different technological success outcomes, defined through the pivot parameters $\boldsymbol{\alpha}$, impact emission abatement costs. For a given portfolio of projects, it is possible to calculate the probability of each success outcome from the elicitation data in Appendix S3 through multiplication of individual project success probabilities. These probabilities, however, are dependent on technology investment decisions, as they will change based on the project portfolio selected. This is the endogenous uncertainty described in Section 1.1, where the probability of a specific outcome is not fixed but rather changes with different decisions. However, such decision dependent probability distributions are typically not amenable for direct use in stochastic optimization, specifically in stochastic programming. To overcome this, we develop a procedure to derive returns functions with fixed probabilities, where each function represents how the values of $\boldsymbol{\alpha}$, i.e. returns from technology investments, change as a function of the investment amounts $\boldsymbol{\Upsilon}$.

A reduced-form portfolio model. The procedure we develop to derive returns functions with fixed probabilities is built upon some insights obtained from the reduced form portfolio model in Baker and Solak (2011), summarized in Appendix S9. Two key conclusions from this work are directly related to our analysis in this paper. First, the authors note that for a given budget level the composition of the optimal portfolio of projects is robust to risk in climate damages.

This conclusion implies that we do not have to explicitly model the individual technology projects in our analysis. Instead, we consider a discrete set of potential R&D investment levels and identify the optimal set of technology projects for each of these investment levels. Recall that for a given portfolio of projects, it is possible to calculate the probability of each possible α value using the elicitation data in Appendix S3. Hence, this process results in a probability distribution for the pivot (and shift) parameters for each investment level. As described later in this section, we use these distributions to derive 'returns functions', in which probabilities are fixed, and pivot and shift parameters change as a function of R&D investment.

Additionally, Baker and Solak (2011) conclude that the optimal amount of R&D funding does change with increases in the riskiness of climate damages in a non-monotonic way. Specifically, when an increase in risk is modeled as a mean-preserving-spread that stretches out the tail, it is found that the optimal amount of R&D funding first increases, and then decreases in risk. The reason has to do with the complex interplay between technology investment and long-term emissions abatement. This result supports the motivation for our current model, in which we use a stochastic dynamic optimization model in order to capture the interactions between climate change damages, abatement, and the impact of R&D under multiple policy frameworks. Through inclusion of several factors such as capital accumulation in the economy, carbon concentration in the atmosphere, and the warming of the deep oceans, our current work is able to represent these complex dynamics in order to provide a convincing policy analysis.

Derivation of stochastic characterizations of returns to R&D. For stochastic characterizations of returns to R&D, we develop a probabilistic mapping from investment decisions Υ_i to the technical change variables α_i, which represent the returns to R&D in technology category i, $i = 1, 2$. As discussed above, this is done by initially considering a discrete set of possible investment levels or budgets. Given a budget level, the reduced-form R&D model identifies an

Budget($mil)	52	108	319	729	961	α_1	52	108	319	729	961	α_1
	Estimates						Actual Data					
Probabilities	0.41	0.24	0.17	0.11	0.11	0	0.41	0.24	0.17	0.11	0.10	0
	0.59	0.34	0.40	0.41	0.35	0.319	0.59	0.34	0.40	0.41	0.35	0.319
	0	0	0.01	0.06	0.12	0.346	0	0	0.02	0.06	0.13	0.346
	0	0.42	0.42	0.42	0.42	0.38	0	0.42	0.42	0.42	0.42	0.38

Table 2 Comparison of the estimated and actual probability distribution data over each possible outcome of α_1 for different levels of investment in CCS.

optimal portfolio of projects for that budget level. Each portfolio is associated with a probability distribution over the possible outcomes of α_i for each technology category i. Since there is only one optimal portfolio per budget level, this allows us to associate a probability distribution over the α_i values to each given budget level.

In the right half of Table 2 we show these values for CCS, with a probability distribution in each budget column and the α_1 values in the last column. An important issue is the selection of the budget levels for this discrete representation of R&D investment. We considered a wide range of possible R&D budgets, and chose those that were optimal investments in some instance of the reduced-form model, or resulted in a significant welfare improvement over other R&D budgets that we considered.

As noted previously, however, endogenous probabilities given in the lower half of Table 2 are not amenable to stochastic optimization, rather we need a mapping where the probabilities are fixed and the α_i values change with the investment decisions. We achieve this by deriving a set of random piecewise linear returns functions, which we denote by $\mathcal{A}_i$, for the two technology categories $i = 1, 2$. Each realization of $\mathcal{A}_i$ maps R&D investment levels Υ_i to technology parameters α_i, i.e. $\mathcal{A}_i : \Upsilon_i \to \alpha_i$. The functions $\mathcal{A}_i$ are piecewise linear as they are defined based on a discrete set of investment levels.

To derive the functions $\mathcal{A}_i$, we started by enumerating all possible functions. As an example, one possible combination for CCS corresponds to the case when α_1 is realized as 0 at

Budget($mil)	52	108	319	729	961	Probability
α_1	0	0	0	0	0	0.11
	0	0	0	0.319	0.319	0.06
	0	0	0.319	0.319	0.319	0.07
	0	0.319	0.319	0.319	0.319	0.17
	0.319	0.319	0.319	0.319	0.319	0.05
	0.319	0.319	0.319	0.319	0.346	0.06
	0.319	0.319	0.319	0.346	0.346	0.04
	0.319	0.319	0.346	0.346	0.346	0.01
	0.319	0.38	0.38	0.38	0.38	0.42

Table 3 Piecewise linear returns functions for CCS, where the central columns show values of α_1 for discrete levels of investment. Each row, which corresponds to a realization of the function $\mathcal{A}_1$, is associated with a probability given in the far right column.

all budget levels. This enumeration, however, results in a large number of functions, which is intractable for optimization purposes. Thus, we perform a scenario reduction process and identify a subset of the possible returns functions that provide a good approximation of the actual data-based distributions. First, we eliminated all returns functions that did not exhibit a dependence between funding levels - that is we assume that if a project is successful at a lower budget level, it will be successful at a higher budget level. Then, we followed a process based on the minimization of the standard deviation of the differences between the actual probability distributions and the probability distributions derived from the subset of the functions. Finally, in order to improve the overall match for solar nuclear, we added in two returns functions that did not exhibit dependence.

Table 3 shows the estimated returns functions for CCS, while the corresponding functions for the solar-nuclear technology category are included in Appendix S10. For example, in Table 3 the third row represents a returns function that has $\alpha_1 = \mathcal{A}_1(\Upsilon_1) = 0$ if $\Upsilon_1 < \$319$ million, and $\alpha_1 = \mathcal{A}_1(\Upsilon_1) = 0.319$ if $\Upsilon_1 \geq \$319$ million. The probability that this particular function is realized is 0.07. As noted above, these functions together with their probabilities provide a very good estimate of the actual probability distributions, with an average standard deviation of the errors of 0.02. Table 2 compares the estimated data with the actual data for CCS, showing the estimated probabilities and the actual probabilities of possible α_1 values at each investment level. They are very close, with the differences being less than 1 percentage point in each case. Similar results hold for the approximations of the solar-nuclear returns functions as well, with differences being less than 4 percentage points in each case. Note that the set of α_i values shown in Table 3 and Appendix S10 are based on the values of the parameters α_{ij} for each project j, which are listed in Appendix S7, and are derived as described in Section 3.2.2. While α_1 is directly based on CCS projects, the combined solar/nuclear parameter α_2 is calculated by first identifying $\alpha_{nuclear}$ and α_{solar}, and then setting $\alpha_2 = 1 - (1 - \alpha_{nuclear})(1 - \alpha_{solar})$.

The resulting stochastic returns functions $\mathcal{A}_i$ are used in conjunction with equation (18). The integration of these piecewise linear functions into the model requires the addition of new variables and constraints. The details of this implementation are described in Appendix S11.

4.1.3. Characterization of a Scenario Set The probabilistic characterizations for the two types of model inputs, i.e. the climate change damages and technical change due to R&D, result in three distinct stochastic entities, which we denote through the random vector $(\pi, \mathcal{A}_1, \mathcal{A}_2)$. Given that possible values for the three parameters are discrete and finite, this random vector can take on a finite number of values. Each of these distinct realizations correspond to a scenario $\omega \in \Omega$ as described in Section 2. The probability of occurrence for each scenario p^{ω} is calculated based on the probabilities of individual parameter value realizations. For example, a sample scenario ω' could correspond to realizations involving the first rows in Table 3 and Appendix S10, and a π value of 1. Assuming that the latter can occur with probability of 1 and using the probabilities shown in Tables 3 and Appendix S10, probability of scenario ω' can be calculated as $p^{\omega'} = (1)(0.11)(0.087) = 0.00957$.

4.2. Step 4: Stochastic Programming Implementation

Stochastic programming is a natural approach for our problem as the interactions represented through the DICE model form a complex structure that prevents the problem from being amenable to other methods such as dynamic programming. This is especially the case as the formulation has many decision variables but relatively few stages.

The optimization problem (1)-(5) can be expressed as one single nonlinear programming problem. However, in order to utilize special algorithmic solution procedures, we further represent this formulation by replacing the first stage decision vectors $\boldsymbol{\Upsilon}, \boldsymbol{\mu}_N, \boldsymbol{y}_N, \boldsymbol{\tau}_N, \boldsymbol{x}_N$ by possibly different vectors $\boldsymbol{\Upsilon}^\omega, \boldsymbol{\mu}_N^\omega, \boldsymbol{y}_N^\omega, \boldsymbol{\tau}_N^\omega, \boldsymbol{x}_N^\omega$, similar to the notation used for second stage decisions. Using this, we can define a problem formulation for each scenario, but at the same time, require that the values of these first stage variables do not depend on the realization of random data. This can be achieved by linking the individual scenario problems through a set of constraints, which are referred to as the nonanticipativity constraints. The nonanticipativity constraints ensure that the decisions in all scenarios are the same for the first 50 years, and are defined explicitly by setting the variables to be equal for each scenario for the first 50 years. In our model, these constraints involve the R&D investment, capital stock and period utility decisions for the first 50 years, i.e. Υ_i, k_t, and u_t, respectively. It must be noted here that the set of decision variables at each period in the model involves a large number of other variables as depicted in the base formulation (6)-(13). However, we show in Appendix S8 that nonanticipativity in the three decisions above is sufficient for the overall model. As fewer number of constraints needs to be used, this result allows for a less complex representation of the stochastic programming model. Moreover, this also enables a tractable implementation of the Lagrangian decomposition procedure which we use as our solution methodology.

Given this structure, the overall stochastic programming model can be defined as follows:

$$\max_{\boldsymbol{\Upsilon}^{\Omega},\boldsymbol{\mu}^{\Omega},\boldsymbol{y}^{\Omega},\boldsymbol{\tau}^{\Omega},\boldsymbol{x}^{\Omega}} \sum_{\omega\in\Omega} p^{\omega} \sum_{t} R_t u_t^{\omega} \tag{19}$$

$$\text{s.t. } J_{\psi}^{\omega}(\boldsymbol{\Upsilon}^{\omega},\boldsymbol{\mu}^{\omega},\boldsymbol{y}^{\omega},\boldsymbol{\tau}^{\omega},\boldsymbol{x}^{\omega}) \le b_{\psi}^{\omega} \qquad \forall\psi,\omega \tag{20}$$

$$\Upsilon_i^{\omega} - \sum_{\omega'\in\Omega} p^{\omega'}\Upsilon_i^{\omega'} = 0 \qquad \forall i,\omega \tag{21}$$

$$k_t^{\omega} - \sum_{\omega'\in\Omega} p^{\omega'} k_t^{\omega'} = 0 \qquad u_t^{\omega} - \sum_{\omega'\in\Omega} p^{\omega'} u_t^{\omega'} = 0 \qquad \forall t\le 5,\omega \tag{22}$$

where the superscripts in $\boldsymbol{\Upsilon}^{\Omega}, \boldsymbol{\mu}^{\Omega}, \boldsymbol{y}^{\Omega}, \boldsymbol{\tau}^{\Omega}$, and $\boldsymbol{x}^{\Omega}$ denote that the optimization is performed over all scenarios in the scenario set Ω. As in the DICE model description in Section 3.2.1, $\boldsymbol{x}^{\omega}$ refers to all other variables in each scenario, and the independent decision variables are $\boldsymbol{\Upsilon}^{\omega}, \boldsymbol{\mu}^{\omega}$ and $\boldsymbol{o}^{\omega}$. In this formulation, objective (19) corresponds to the objective function (1) in the general formulation and represents the maximization of expected total social utility over all scenarios. Constraints (20) define the corresponding set of constraints for any given scenario ω. These constraints, each of which is indicated by $\psi = 1, \ldots, \Psi$, involve the standard economic relationships as given by (7), (9), (10), (12), (13), as well as the extended relationships modeled by equations (16), (18) and the required constraints for the piecewise linear mappings described in Appendix S11. The variables in each of these constraints are indexed by a superscript ω, and the constraints are defined separately for each $\omega \in \Omega$. Constraints (21)-(22) are the nonanticipativity constraints ensuring that decisions in the first 50 years are the same for all scenarios. The structure used in the formulation of the nonanticipativity constraints accounts for the scenario probabilities, and prevent the ill-conditioning in the Lagrangian dual as discussed by Louveaux and Schultz (2003). Note that constraints (20)-(22) together define a reformulation of constraints (2)-(5) in the general formulation described in Section 2.

Model (19)-(22) is a stochastic nonlinear programming problem that can be solved through decomposition methods, as the size of the scenario set does not allow for a direct solution. As a solution approach, we use a Lagrangian decomposition based procedure, which is similar to the method described by Caroe and Schultz (1999). The details of this implementation, as well as some additional computational improvement procedures are described in Appendix S12.

5. Experimental Setup for R&D Policy Analysis

In this section we discuss the policy experiments we ran with our framework. We start by briefly discussing some assumptions about the opportunity cost of investment. We then describe a number of alternative policy environments and the different risk cases we consider.

5.1. Opportunity Cost

The R&D funding levels used in the elicitations and reported in the tables in Appendix S3 represent the amount of money going into the hands of high quality researchers in the appropriate areas; this does not account for additional costs to society. In order to address this issue, our baseline assumption is that the opportunity cost of investing in R&D is 4 times the out-of-pocket cost (i.e. $\kappa = 4$ in equation (16) for the base case). This assumption reflects the current state of the literature (Nordhaus 2002, Pizer and Popp 2008), but in fact there is very little research directed at determining what this opportunity cost actually is. Thus, we perform sensitivity analysis over the parameter κ, and discuss it in our analysis in Section 6.

5.2. Alternative Policy Environments

Following Nordhaus (2008), we consider a number of different policy environments which prescribe different alternative strategies in dealing with climate change. Specifically, we consider six policy environments which we refer to as DICE Optimal, Stern, Stern Fixed, Gore, Kyoto Strong, and Lim 2, as well as a baseline no-controls case. We choose these policies because they are representative of the range of policy recommendations being debated around the world. Here we describe these policies, which are also summarized in Table 4. For each

Policy	Abatement	Key Characteristics
Baseline	no controls	-
DICE Optimal	optimal	-
Stern	optimal	abatement chosen under low interest rate
Stern Fixed	optimal	abatement and R&D chosen under low interest rate
Gore	lower bound btwn 0.25-0.95	limited participation
Kyoto Strong	fixed for 150 yrs	limited participation
Lim 2	optimal	upper bound on temperature

Table 4 Attributes of policies considered.

policy environment we assume there is no knowledge of technological success and damages until year 2055, and we run the model out for 400 years.

In the "DICE Optimal" policy the model chooses the optimal R&D investment and abatement path. The "Baseline" case models the levels of major economic and environmental variables as they would occur without any climate-change policies, i.e. abatement is forced to be 0 in all periods after the first. The "Stern" policy is intended to reflect the policy suggestions laid out in the Stern Report (Stern 2007). Nordhaus (2008) identified the key difference between Stern and DICE as being the very low discount rate in the former. Thus, this policy is implemented by first running the DICE model at a very low discount rate. We then take the resulting abatement levels and fix those in the model with the default discount rate. This is so the results of all the policies can be evaluated at one common interest rate. For our implementation, we have run two versions of this policy. In one case (referred to as "Stern Fixed"), R&D investment is fixed as calculated from the run with the very low discount rate. In the other case (referred to as "Stern"), R&D investment is chosen in the second run based on the DICE discount rate[4].

The "Gore" policy is intended to reflect the policy suggestions laid out by Al Gore (Gore 2007). This policy fixes a lower bound for abatement of $0.25, 0.45, 0.65, 0.85$ for the periods beginning 2015, 2025, 2035, and 2045 respectively. Thereafter the lower bound for abatement is fixed at 0.95. However, for our model, we assumed that when climate change uncertainty

[4] The reader can refer to Nordhaus (2008) for more information on how the interest rates are modeled.

	No risk (1)	Medium risk (2)		High risk (3)		Very high risk (4)		Intermediate (5)		
GDP Loss	1.1%	0.0%	3.3%	0.0%	20.0%	0.0%	40.0%	0.0%	1.1%	20.0%
Probability	1.000	0.667	0.333	0.945	0.055	0.973	0.028	0.309	0.673	0.018
π	0.003	0.000	0.009	0.000	0.063	0.000	0.167	0.000	0.003	0.063

Table 5 Probability distributions defining climate change damage uncertainty.

is realized with a no damage outcome, abatement is chosen optimally. The Gore policy also reflects limited participation in the early periods, in which not all countries and regions will participate in the abatement. Specifically, it is assumed that the participation rate will increase gradually from 0.6 to 1 over the next 50 years.

“Kyoto Strong” represents a very aggressive, but potentially achievable, international agreement on climate change. This is intended to follow the spirit of the Kyoto Protocol, but to continue on indefinitely and have more and more nations join on through time. In this policy, abatement is fixed for the first 150 years. To be able to model the learning about climate damages and R&D, we altered this aspect, as in the Gore policy, by allowing abatement to be chosen optimally in the case where damages are zero. Thus, abatement does not respond to the outcome of R&D for the first 150 years, nor to higher than expected damages. Also, similar to Gore, the cost of abatement is increased at earlier stages when fewer countries have joined in. After 150 years, future abatement is chosen optimally and responds to the particular scenario.

Finally, the “Lim 2” policy adds a single constraint that limits the average global temperature increase to 2°C. In our modeling framework this constraint is only minimally binding.

5.3. Risk Cases

One of our central questions is how uncertainty about climate change damages impacts near-term investments. In order to address this question, we consider multiple cases for uncertainty over climate damages. Table 5 shows the five cases we consider. Each probability distribution is given a name in the top row. The second row shows the percent GDP loss given a 2°C increase in global mean temperature as calculated using equation (11). We choose this value as our anchor because it is used to calibrate the DICE model, and is the value used in the elicitations in Nordhaus (1994). We use mean-preserving spreads around this value. That is, each probability distribution has a mean GDP loss of 1.1% given a 2°C warming. The second row shows the probabilities of each outcome in that distribution. The last row shows the value of the parameter π for each respective outcome. While previous work uses mean-preserving spreads around π, the damages in the DICE model are concave in π, therefore we hold the mean of the expected GDP loss constant in the mean-preserving spreads used in this paper.

6. Results and R&D Policy Analysis

In this section, we discuss the results from the optimization model, focusing on the impact of R&D on societal costs and on the range of scenarios considered in the analysis. While we describe several conclusions derived from our model and data, we emphasize that the analysis boundaries are tight, and therefore the results may not be generalizable.

6.1. Optimal Investment in R&D

We find that *the optimal investment in the R&D projects considered is quite robust, both across different policy environments and across different risk cases.* Figure 1 illustrates the optimal investment across risk and policy environments. The horizontal axis represents risk cases 1 – 3 from Table 5. Since the optimal investment is the same in each case for DICE Optimal, Gore, and Kyoto Strong, we have graphed these together; the same goes for Stern and Lim 2. We see that in 12 out of the 15 cases we are showing, the total optimal investment is equal to $\$5,303$ million[5]. In the three other cases, the optimal investments drop by less than $300 million, a relatively small amount. This robustness can be partly explained by referring to the reduced-form R&D model where the portfolio of technology projects was also observed to be robust to

[5] We show the allocation of the total investment in research areas for the optimal levels of investment in Appendix S13.

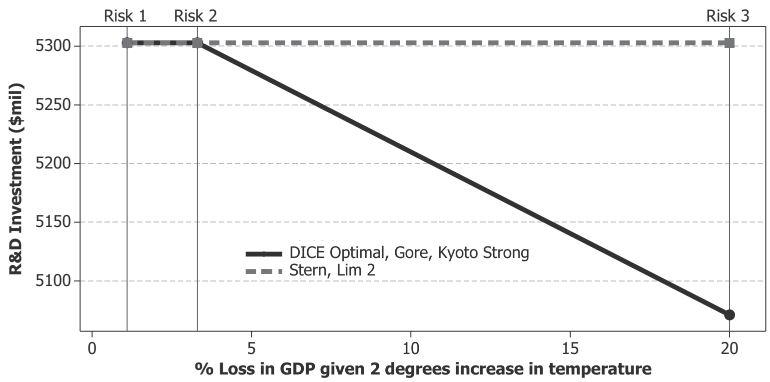

Figure 1 **Optimal investment across risk and policy scenarios, where the horizontal axis values correspond to the GDP loss for the high damage outcome of risk cases 1 – 3 from Table 5.**

climate damages for any given R&D budget. This is because the elicitation data revealed the technologies to be quite diverse, with some projects clearly superior to others. Here, however, our robustness result is even stronger. We find that the same level of funding is optimal over a wide range of very different abatement paths and different risk levels.

Not shown in Figure 1 is the DICE Optimal investment under very high risk, which falls to $\$2,696$ million. In general, we see a somewhat monotonic response to risk in our results, with the optimal investment in R&D decreasing in risk under at least some of the policies. This is because when damages are very high, abatement is at 100% with or without technical change. A mean-preserving spread that increases the magnitude of the damages will simultaneously reduce the probability of those damages. Thus, an increase in risk of this kind leads to a lower probability of full abatement, and therefore a lower expected value for technical change.

The optimal investment is also fairly robust to assumptions about the opportunity cost of R&D investments. If the opportunity cost multiplier κ is between $1-4$, then the optimal investment is stable as above at $5,303 million. If it is between $5-6$, then the optimal investment is slightly lower at $5,071 million. If the opportunity cost is between $7-10$, the optimal investment drops to $2,696 million. Thus, the investment in the R&D projects considered is not very sensitive to assumptions about opportunity cost until the opportunity cost gets very high.

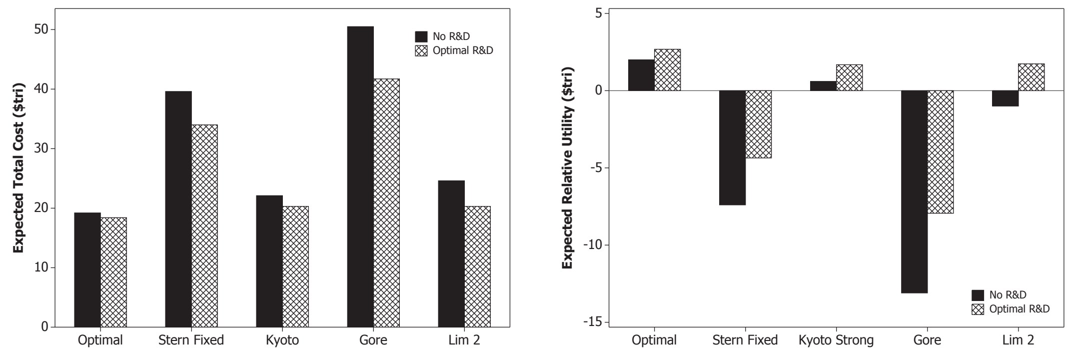

(a) Effect of R&D on total costs (b) Expected relative utility of policy interventions

Figure 2 Expected costs and utilities of policy interventions

The optimal investment is much higher in Stern Fixed, in which the investment is chosen with a very low discount rate. The optimal action in this policy environment is to invest $\$21,132$ million, the maximum amount we have available in our model. This is not surprising, and underlines the importance of coming to an agreement on discount rates.

6.2. Impact of R&D on Expected Policy Costs

In Figure 2, we illustrate the impact of R&D on the different policy environments. Figure 2(a) shows the expected total cost of each policy with and without R&D, while Figure 2(b) shows the expected relative utility of each policy intervention with respect to the Baseline.

The vertical axis in Figure 2(a) represents the expected net present value of the cost of abatement plus the cost of climate damages in trillions of 2005 dollars. The vertical axis in Figure 2(b) represents expected utility in the same units. Each bar represents the extra utility gained (or lost) from the Baseline by implementing the policy intervention. The darker bars are in the absence of R&D and are very similar to the results in Nordhaus (2008). The lighter bars are when R&D is available and chosen optimally. Note that some of the policy environments are improvements over doing nothing, whereas some are worse than doing nothing, at least as evaluated within our framework. The Stern Fixed policy is too stringent in the DICE model, as it is chosen in response to a very low interest rate, but evaluated under a higher interest rate.

The first result here is that *the availability of R&D is more valuable in the second best policy environments*. The value of having R&D is greater in the non-optimal environments and is greatest in the Gore environment. Of particular interest are the Kyoto Strong and Lim 2 results. Kyoto Strong is a possibly implementable policy. In the absence of R&D, it is barely better than doing nothing. However, with R&D it becomes clearly positive, almost equivalent to the optimal without R&D. The Lim 2 goes from being a net loss to a net benefit with R&D.

Figure 3(a) shows the expected utility of policy intervention for the DICE Optimal policy, comparing no R&D, optimal R&D (of between 5.0 - 5.3 billion), and full R&D investment in all of the technologies (of $\$21$ billion), for three risk cases. What is striking here is the asymmetrical effect of over-investment relative to under-investment: over-investment has a smaller downside. To further analyze this, we look at the expected utility of policy intervention for the DICE-optimal policy for R&D investments that are marginally higher and lower than the optimal. Specifically, we considered the four following budgets (in millions): $5,689, $5,303, $5,071, and $4,743. The middle two values are the two budgets that are optimal in Figure 1. The higher and lower numbers increase and decrease these central values by $386 million and $328 million, respectively. Figure 3(b) shows that while the two central budget levels lead to very similar expected utility, *there is an asymmetric effect of increasing or decreasing the budget further, with under-investment being more costly than over-investment*.

6.3. Impact of R&D on the Range of Scenarios

Figure 4(a) shows the range of abatement paths over time in the Risk 1 case for DICE Optimal when R&D is invested optimally. There are a total of 36 scenarios, each depending on the success outcomes of the technologies. On the right edge of the graph we show the probability of being in a group of scenarios. The first cluster, with probability of 45%, is associated with scenarios in which there is success in both nuclear and CCS. The next cluster, with probability

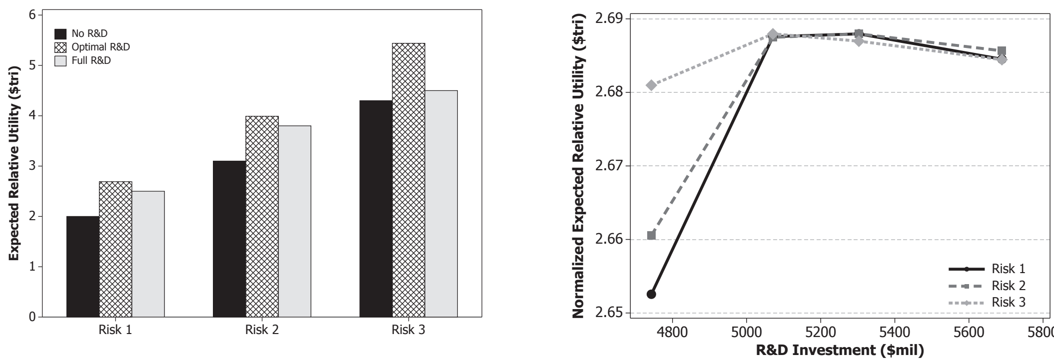

(a) Expected utility of policy intervention for the DICE Optimal policy for no and full R&D budgets

(b) Expected utility for the DICE Optimal policy for R&D budgets marginally different from optimal

Figure 3 Expected utility of policy intervention.

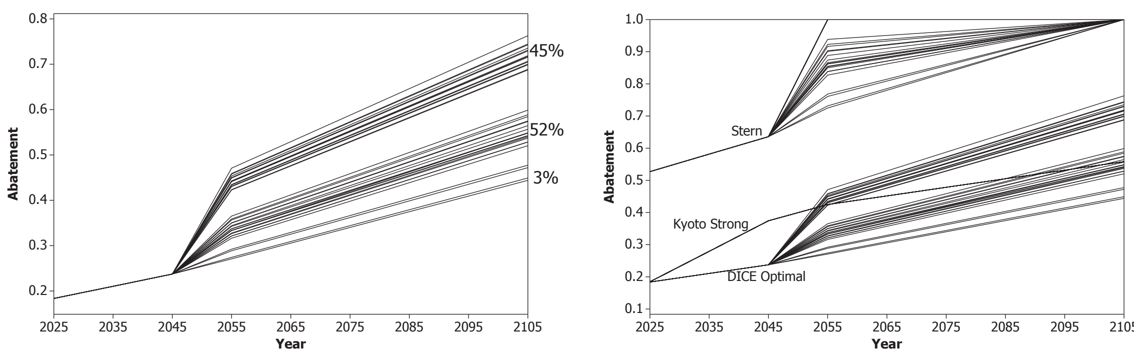

(a) Range of abatement paths in the Risk 1 case for DICE Optimal

(b) Comparison of the range of abatement paths from DICE Optimal, Stern, and Kyoto Strong

Figure 4 Ranges of abatement paths for different realizations of technical uncertainties.

of 52%, includes scenarios in which either nuclear or CCS fails. The lowest line is the case where all technology fails. In Figures 4(b) and 5 we just show the range of paths without associated probabilities, since they follow a similar pattern. Figure 4(b) compares the range of emissions paths from the DICE Optimal, Stern, and Kyoto Strong policies. Note that in the absence of technical change the Kyoto Strong path is higher than the optimal abatement path. With technology, however, it falls about in the middle of the optimal paths. Thus, *the presence of R&D greatly enhances the value of the fixed emissions path prescribed by Kyoto Strong.*

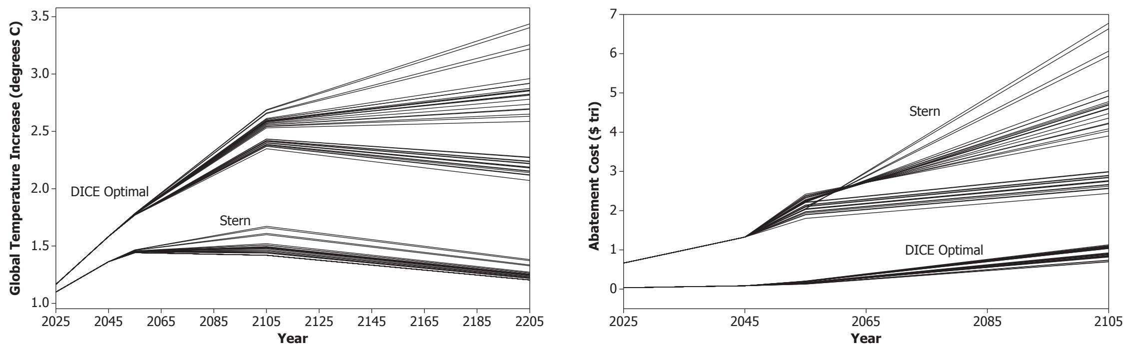

(a) Range of temperature paths in the Risk 1 case for DICE Optimal and for Stern

(b) Range of abatement cost paths in the Risk 1 case for DICE Optimal and Stern

Figure 5 Ranges of temperature and abatement cost paths.

Figure 5(a) shows the range of temperature paths, i.e. the change in the average global temperature over time, in the Risk 1 case for DICE Optimal and for Stern. From this figure we can make three observations. First, all the DICE Optimal paths are above 2°C between 2075 and 2200, while all Stern paths are always below. What we can conclude is that *Stern with no advances in technology will lead to lower temperatures than DICE Optimal with great advances in technology*. Second, the impact of R&D on the temperature is much stronger in the DICE Optimal policy than in the Stern policy. As the third observation, we note that all Stern paths and the DICE Optimal paths with the most successful R&D peak in temperature between 2100 and 2200. Temperature will peak in any scenario after abatement hits 100%. All paths hit full abatement eventually, but R&D can significantly affect the timing.

Figure 5(b) shows the abatement costs for all scenarios for the Risk 1 case of DICE Optimal and Stern. We see that until 2105 R&D has a much larger impact on Stern than on DICE Optimal. In fact, Figure 5(b) is almost the opposite of Figure 5(a) showing the temperature paths. *If we are in a policy environment in which abatement is relatively high, then R&D will have a large effect on abatement costs and a smaller effect on emissions, temperature and other physical variables*. If, on the other hand, we are in a policy environment that leans toward lower abatement, then R&D will have a large effect on emissions and temperatures, and a smaller effect on costs. The robustness of R&D investment to different policy environments and risk cases can be partially explained by this phenomenon. Even though the policy environments are radically different in their abatement paths, technological change has a role to play in both: cost reduction when abatement is high, and improved environmental impacts resulting from higher abatement when abatement is generally lower.

As an additional analysis, we also present in Appendix S14 how investment into R&D impacts the riskiness of the policy outcomes, where we conclude that in general R&D provides risk reduction.

7. Conclusions and Further Policy Insights

Finding an optimal R&D portfolio in the face of climate change is a challenging problem, and the available data is sparse and not as airtight as we would like. It is, however, a real and pressing problem faced by the U.S. and other governments around the world. Thus, our approach is to use the best data available and explicitly include uncertainty in our analysis, in order to arrive at robust insights *conditional on the current state of knowledge*. With this aim in mind, we have developed multiple models and implemented data-based uncertainty on the returns from R&D into a newly developed stochastic version of an IAM in order to get insights about the optimal technology R&D portfolio in the face of climate change. Overall, we developed a general framework to determine optimal R&D portfolios through a dynamic stochastic model. Thus, this work provides a framework which can be updated as new information becomes available, such as more detailed elicitations about these technologies or data on other technologies.

First, we have found that, given our data based on expert elicitations, focused investments in Nuclear LWR and HTR, as well as in CCS and solar technologies are very robust. The optimal investment in these projects was robust to the policy environment, to the riskiness of climate damages, and to opportunity cost. Conversely, not investing in the Nuclear FR and Solar 3rd G projects was also robust. We do see a lower investment when there is a very small probability of very high damages; and a larger investment when the interest rate is very low; but otherwise the optimal investment falls within a very small range. This robustness is interesting, as for example the Gore and the DICE Optimal policies are different in most other ways. It is also a very useful result, implying that near-term decisions to invest in the projects considered in this analysis may not depend heavily on the outcomes of long-term climate policy decisions.

This robustness is driven by two effects. The first is a result of using *data* (rather than theoretical explorations): we found that individual projects were quite differentiated, with some projects having relatively lower costs, large impacts if successful, and high probabilities of success, while other projects did not perform well on all or some of these aspects. Clearly, our conclusions are conditional on the boundaries of the elicitation data, as robust results can not be guaranteed for any set of data. However, we believe that robustness is more likely than not when using real data, since there will always be a relatively narrow band of benefits-to-costs that will put a project on the knife's edge.

The second driver relates to the *role of R&D* in the different scenarios. This role varies considerably, both with the policy environment and with the uncertainty over damages. In policy environments in which abatement is fixed or tends to be very high (near 1), R&D primarily has a 'cost-side benefit': the environmental variables are less affected while the cost of abatement is significantly affected. This group of policies and risk cases include the Stern and Gore policies and also high risk cases. On the other hand, in instances in which abatement is relatively low in the absence of R&D, R&D primarily has an 'environmental-side benefit': the environmental variables are significantly affected, while the cost of abatement has only small effects, and in fact sometimes is higher given a much higher level of abatement. These two very different roles mean that technological change ends up having an important role to play whether abatement is high or low. This insight would be unlikely to arise outside of our framework that combines a dynamic optimization model with data-based probability distributions.

Second, we have shown that a larger-than-optimal investment in technology is less costly than a smaller-than-optimal investment. Thus, it appears that policy makers should prefer to err on the high side rather than the low side of R&D investment. Given this result, the level of robustness, and the deep uncertainty about climate damages, our observations support a conclusion that investing roughly $5 billion in these technologies probably makes sense.

Our research is plagued by the same difficulties that plague all climate change research – the optimal investment depends on the interest rate used to value the far future. We see that the optimal investment in R&D is considerably higher – in fact full funding in all the projects we considered – when evaluated at the low Stern interest rate. If policy makers believe that the 'appropriate' interest rate is no higher than that in Nordhaus (2008) (since very few economists are making that argument), this again suggests that policy makers err on the side of higher investments rather than lower. Assumptions about the opportunity cost of R&D investments have little impact on the results in this study.

There are some important caveats to these conclusions that need to be considered in a final investment decision. First, this is a 'lumpy' problem, with the projects defined by discrete investment levels, some of which are much higher than others. Moreover, the analysis is based on a specific set of projects deemed relevant for R&D portfolio optimization by a set of experts. Hence, biases inherent in any expert elicitation process may exist in the set of projects considered. This is not atypical of R&D portfolio problems, and is driven by the difficulty of assessing potential R&D projects. Nevertheless, some of the robustness may be driven by this characteristic. Future work may be aimed at minimizing this problem. More generally, data gathered from expert elicitations are typically subject to a number of biases and may vary greatly depending on the structure of the elicitation process. Certainly, given the scale of the climate change problem, more and better information on the potential of energy technology R&D is likely of great value (Baker and Peng 2012).

In addition, the models we worked with (and we believe this is true for all models) could not account for the socio-political aspects of nuclear energy. In particular, concerns about proliferation are not adequately reflected in this analysis. Thus, nuclear may be a riskier investment than we show. Also, there is only a weak understanding of how intermittent renewables, such as solar photovoltaics, will be able to be integrated into the grid on a large scale. Thus, while the models we used consider this problem in a reasonable way, it is quite possible that the impact of improvements in solar photovoltaics will be larger than the current models show, especially if simultaneous investments are made in the grid and grid integration. Thus, these two technology-specific aspects should be considered in a final portfolio allocation. More generally, we made a number of assumptions along the way in order to integrate the data and the many components of the framework. Hence, as is typical in approaches to such complex problems, the results should be interpreted with caution and may not be generalizable.

Beyond the specific contributions to climate change energy R&D policy, this paper provides an example of a framework for combining elicitation-based probabilistic data on future uncertain systems and multiple economic models into a tractable stochastic decision framework. We introduced the idea of random return-to-R&D functions, which were then effectively integrated into a novel representation of a highly nonlinear stochastic problem. Overall, we were able to integrate probabilistic data into a fully dynamic model in order to derive robust policy insights. This framework may be applied not only to the broad and important field of energy technology portfolio selection, but also to other public policy areas such as R&D into space exploration, health, and military, as well as agencies such as the Environmental Protection Agency who face choices of a portfolio of policies that have uncertain uptake and response on firm side, and uncertain benefits (in the sense of poorly understood pollutants).

Acknowledgments

This research was partially supported by NSF under award number SES-0745161, and the Precourt Energy Efficiency Center at Stanford University.

References

Anadon, L., M. Bunn, G. Chan, M. Chan, C. Jones, R. Kempener, A. Lee, N. Logar, V. Narayanamurti. 2011. Transforming U.S. energy innovation. Tech. rep., Harvard Kennedy School.

Baker, E. 2009. Optimal policy under uncertainty and learning in climate change. *Journal of Public Economic Theory* **11**(5) 721–747.

Baker, E., H. Chon, J. Keisler. 2008. Nuclear power. Working Paper, UMass Amherst.

Baker, E., H. Chon, J. Keisler. 2009a. Advanced solar R&D. *Energy Economics* **31** S37–S49.

Baker, E., H. Chon, J. Keisler. 2009b. Carbon capture and storage. *Climatic Change* **96**(3) 379–408.

Baker, E., Y. Peng. 2012. The value of better information on technology R&D programs in response to climate change. *Environmental Modeling and Assessment* **17**(1) 107–121.

Baker, E., E. Shittu. 2008. Uncertainty and endogenous technical change in climate policy models. *Energy Economics* **30**(6) 2817–2828.

Baker, E., S. Solak. 2011. Climate change and optimal energy technology R&D policy. *European Journal of Operational Research* **213**(2) 442–454.

Blanford, G.J. 2009. R&D investment strategy for climate change: A numerical study. *Energy Economics* **31** S27–S36.

Brenkert, A.S., S. Smith, S. Kim, H. Pitcher. 2003. Model documentation for the MiniCAM. Tech. Rep. PNNL-14337, Pacific Northwest National Laboratory.

Caroe, C.C., R. Schultz. 1999. Dual decomposition in stochastic integer programming. *OR Letters* **24**(1) 37–45.

Clarke, L., J. Edmonds, V. Krey, R. Richels, S. Rose, M. Tavoni. 2009. International climate policy architectures. *Energy Economics* **31** S64–S81.

Clarke, L., P. Kyle, M.A. Wise, K. Calvin, J.A. Edmonds, S.H. Kim, M. Placet, S. Smith. 2008. CO_2 emissions mitigation and technological advance. Tech. rep., Pacific NW Nat. Lab.

Edmonds, J.A., J.F. Clarke, J.J. Dooley, S.H. Kim, S.J. Smith. 2004. Stabilization of CO_2 in a B2 world. *Energy Economics* **26**(4) 517–537.

Gore, A. 2007. Moving beyond Kyoto. New York Times, July 1.

Hora, S. 2004. Probability judgments for continuous quantities: linear combinations and calibration. *Management Science* **50**(5) 597 – 604.

InterAcademy Council. 2010. Climate change assessments: Review of the processes & procedures of the IPCC. Tech. rep., InterAcademy Council.

Lewis, N.S., D.G. Nocera. 2006. Powering the planet: Chemical challenges in solar energy utilization. *Proceedings of the National Academy of Sciences* **43** 15729 – 15735.

Louveaux, F.V., R. Schultz. 2003. *Handbooks in Operations Research and Management Science*, vol. 10, chap. 4. Elsevier Science, 213–266.

National Research Council. 2007. *Prospective Evaluation of Applied Energy Research and Development at DOE (Phase Two)*. The National Academies Press, Washington D.C.

Nordhaus, W.D. 1993. Rolling the DICE. *Resource and Energy Economics* **15**(1) 27–50.

Nordhaus, W.D. 1994. *Managing the Global Commons*. MIT Press, Boston, MA.

Nordhaus, W.D. 2002. Modeling induced innovation in climate change policy. A. Grubler, N. Nakicenovic, W. D. Nordhaus, eds., *Technological Change and the Environment*. RFF.

Nordhaus, W.D. 2008. *A Question of Balance*. Yale University Press, New Haven, CT.

Nordhaus, W.D. 2011. Designing a friendly space for technological change to slow global warming. *Energy Economics* **33**(4) 665 – 673.

Pizer, W. A., D. Popp. 2008. Endogenizing technological change. *Energy Economics* **30**(6) 2754–2770.

Stern, N. 2007. *The Economics of Climate Change*. Cambridge Univ. Press, Cambridge, MA.

Tversky, A., D. Kahneman. 1974. Judgement under uncertainty. *Science* **185** 1124–1131.

Supporting Information

Additional supporting information may be found in the online version of this article:

Appendix S1. Definitions of Acronyms

Appendix S2. Definitions of Variables and Parameters

Appendix S3. Summary of Expert Elicitation Results

Appendix S4. Description of Inapplicability of Other Elicitation Data

Appendix S5. Description of Cost of Abatement in DICE

Appendix S6. Representative Marginal Abatement Cost Curves

Appendix S7. Pivot Parameter Values for Individual Technology Projects

Appendix S8. Proofs of Analytical Results

Appendix S9. Reduced Form R&D Model

Appendix S10. Returns Functions for the Solar-Nuclear Technology Category

Appendix S11. Representation of Stochastic Returns Functions

Appendix S12. Description of the Solution Procedure

Appendix S13. Allocation of Total Investment under Different Optimal Investment Values

Appendix S14. R&D and Riskiness of Outcomes