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

(Authors' names blinded for peer review)

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 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

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 models that use 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?

In order to address these questions and provide policy insights, we develop a multi-step multi-model framework 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 relevant research in the area, and discuss how our analysis and major findings contribute to the existing literature.

1.1. General Decision Framework

There are two key near-term avenues for societal response to climate change. The first and most direct avenue 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 avenue would be the determination of the optimal path of emissions in future years, emissions allocations, or the level of a carbon tax. A second avenue of response is to invest in *energy technology R&D* so that emissions abatement will be less costly in the future. This second avenue is the main focus of this paper.

These two avenues of response to climate change, i.e. abatement and energy technology R&D, are not independent. 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. To address this, 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. For the near term, the decision makers face potential investment decisions over a set of alternative energy technology R&D projects with well-estimated costs but uncertain returns, as well as decisions on how much to reduce emissions given current abatement costs. Long term decisions, on the other hand, correspond to future emission reductions under information to be learned on impacts of climate change and on future abatement costs, where the latter will depend on the level of success in invested technology projects.

We illustrate this decision process through the influence diagram in Figure 1 as follows. The near term decisions (with decisions represented by rectangular nodes) are how much to invest in which technologies and the level of short term abatement, where abatement is defined as the fraction of

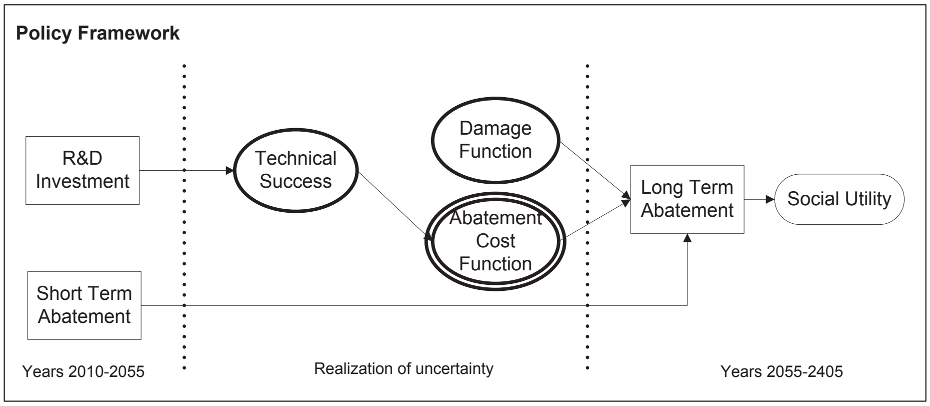

Figure 1 **Influence diagram for the decision process to maximize social utility in the presence of climate change and technological uncertainty.**

emissions reduced below business-as-usual level. The arrow from the R&D investment decision node to the oval uncertainty node for technical success means that the probability distribution of technological success depends on the projects that are funded, 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. The next relationship is a deterministic mapping between technical success and abatement costs, which implies that technological success will impact the costs of reducing emissions. The second stochastic characterization in Figure 1 (also denoted by an oval uncertainty node) 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, referred to as long term abatement in Figure 1, will be determined after information about future abatement costs and the damages becomes available. The objective for this 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 limitations 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 theoretical 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). Two existing R&D studies, namely Blanford and Weyant (2007) and Blanford (2009), use illustrative probability distributions and simply assume there are decreasing returns to scale in R&D investment. Our study differs from these papers 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. Using those insights and other inputs, we develop a complex multi-period multi-constraint R&D optimization framework that is able to represent the long-term dynamic interactions between R&D and climate change, and provide a comparative analysis over several potentially implementable policy environments, such as those proposed by Al Gore (Gore 2007), the Stern report (Stern 2007), and the Kyoto Protocol.

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. These models have been used to perform extensive sensitivity analysis, generally looking at the impact on the present value of the long term costs of abatement for different environmental goals and different assumptions about technologies (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. (2011b), where the authors combine empirical data with a detailed IAM-based analysis to perform portfolio optimization. Unlike our study, however, 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.

1.3. Contributions and Major Findings

Given the existing studies on energy technology R&D portfolio analysis, the specific contributions of this paper to the literature can be summarized as follows:

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

(1) We build on the technological detail of an IAM, and go beyond sensitivity analysis to specifically incorporate empirically-based uncertainty into climate change energy technology policy analysis. To the best of our knowledge, our analysis is the first such integrated approach that combines elicited probabilistic data and stochastic optimization with detailed assessment models in a dynamic framework in order to support R&D investment decisions.

(2) We use this novel stochastic model to analyze the impacts of uncertainty and learning about technology and climate damages on the interplay between R&D, technological change, and optimal emissions policy within multiple policy frameworks. Thus, we provide a study that goes beyond theoretical analysis on the one hand, and data-based sensitivity analysis on the other hand.

(3) While we focus on a specific application, our methodology – applying elicitations to complex economic decisions, and combining multiple models to increase tractability in large problems – may be useful in a number of applications.

(4) Using our framework and the best available data representing the current state of knowledge, we are able to reach various policy conclusions on the energy technology R&D portfolio problem. Our major findings show 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 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.

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]. Similarly, the vectors $\boldsymbol{x}_N$ and $\boldsymbol{x}_L$ are used to represent the set of other decision variables in each stage that define the relationships between climate change, economy and social utility. 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 can be expressed in general form as follows:

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

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

$$\text{s.t. } \boldsymbol{G}_N(\boldsymbol{\Upsilon},\boldsymbol{\mu}_N,\boldsymbol{x}_N)\le \boldsymbol{b}_N \tag{2}$$

$$\boldsymbol{G}_L(\boldsymbol{\Upsilon},\boldsymbol{\mu}_N,\boldsymbol{x}_N,\boldsymbol{\mu}_L^\omega,\boldsymbol{x}_L^\omega,\omega)\le \boldsymbol{b}_L(\omega) \qquad \forall\omega \tag{3}$$

where the superscript in $\boldsymbol{\mu}_L^\omega$ and $\boldsymbol{x}_L^\omega$ denote the dependence of these decision variables on the realized values of the stochastic parameters, while the superscript in $\boldsymbol{\mu}_L^\Omega$ and $\boldsymbol{x}_L^\Omega$ implies that the optimization is performed over the corresponding variables in all scenarios, as we let $\boldsymbol{\mu}_L^\Omega=\{\boldsymbol{\mu}_L^1,\ldots,\boldsymbol{\mu}_L^{|\Omega|}\}$ and $\boldsymbol{x}_L^\Omega=\{\boldsymbol{x}_L^1,\ldots,\boldsymbol{x}_L^{|\Omega|}\}$.

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, the objective 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. Similarly, the constraint sets (2) and (3) correspond to the relationships defining the interplay between climate change, economy and social utility for the near-term and long-term decision problems, respectively. Note that the second stage constraint set (3), and thus the optimal long-term abatement policy, is dependent on the technology investments $\boldsymbol{\Upsilon}$, near-term abatement $\boldsymbol{\mu}_N$, and other related decisions $\boldsymbol{x}_N$ made in the initial decision stage.

While a general mathematical model description is possible as described above, there are several challenges that need to be overcome for 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 $U_s(\cdot)$ and the constraint sets $\boldsymbol{G}_s(\cdot)$, $s=N,L$: 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)-(3): 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 each of these issues to identify policy results for energy technology R&D under climate change. The first two steps 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 the component development process separately for each of the first two steps listed in Section 2 above. For some parts of these processes we utilize results from significant modeling and analysis efforts, which 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. A key part of our multi-step multi-model approach, which is directly related to the definition of the investment decision vector $\boldsymbol{\Upsilon}$, was to determine energy technologies for consideration and get expert opinions from scientists and engineers familiar with these technologies. 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 corresponding research areas for each technology are listed in Appendix S3. These projects were chosen in conjunction 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. To define project return characteristics, we assume that each project can be invested in at one of multiple potential levels, where investment is 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 return characteristics for each project are based on the probability of success in that project, which is defined through expert elicitations summarized in Appendix S3. These elicitations are described in detail by Baker et al. (2008), Baker et al. (2009a) and Baker et al. (2009b). 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). The specific probabilities of success for different investment levels of each project reflect an aggregation of the individual experts' judgments. Expert input was used to define both the funding levels and the endpoints to be assessed.

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, the 2010 InterAcademy Council review of the climate change assessment of the Intergovernmental Panel on Climate Change specifically suggested 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 note here that while other elicitation data exists on the three technologies we consider, they are not applicable to the type of R&D portfolio analysis studied in this paper. For example, in some of the other studies the subtechnologies or projects in each technology are not differentiated at all (National Research Council 2007, Bosetti et al. 2012), while in one study only one project is considered (Rao et al. 2006). In a few other existing elicitations, each expert evaluated the project they thought was the most promising project (Curtright et al. 2008, Anadon et al. 2011a, Chan et al. 2011). Such data would not work very well for a general R&D analysis, as it would result in an increased level of bias in the models due to not aggregating over multiple experts. Hence, the expert elicitation data used in our analysis represents the most appropriate currently available data for energy technology R&D portfolio policy analysis.

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

The modeling of the social utility functions $U_s(\cdot)$ and the constraint sets $\boldsymbol{G}_s(\cdot)$ for $s = N, L$ involves multiple phases. First, the complex relationships between climate change, emissions abatement, economy and social utility are represented in basic form through the utilization of 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 then 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 formulation of the DICE model is given in Nordhaus (2008), and is also summarized through equations (4)-(11) below. The summary formulation utilizes the following parameters and variables:

Parameters

- R_t : utility discount factor for period t
- A_t : level of total factor productivity in period t
- S_t : ratio of uncontrolled emissions to output in period t
- E_t : emissions from deforestation in period t
- L_t : population and labor input in period t
- β : elasticity of marginal utility of consumption
- γ : elasticity of output with respect to capital
- σ : rate of depreciation of capital

Variables

- o_t : consumption of goods/services in period t
- y_t : net output of goods/services in period t
- y_t^g : unadjusted output in period t
- k_t : capital stock in period t
- l_t : investment in period t
- e_t : total carbon emissions in period t
- μ_t : emissions abatement in period t
- τ_t : atmospheric temperature in period t
- u_t : social utility in period t

$$\max_{\boldsymbol{\mu},\boldsymbol{x}} \quad \sum_t R_t u_t \tag{4}$$

$$\text{s.t.} \quad u_t = L_t \frac{(\frac{o_t}{L_t})^{1-\beta} - 1}{1-\beta} \qquad \forall t \tag{5}$$

$$y_t = o_t + l_t \qquad \forall t \tag{6}$$

$$k_t = l_{t-1} + (1-\sigma)k_{t-1} \qquad \forall t \tag{7}$$

$$\tau_t = H\left(\tau_{t-1}, e_t\right) \qquad \forall t \tag{8}$$

$$y_t = \frac{1 - c_D\left(\mu_t\right)}{D_D\left(\tau_t\right)} y_t^g \qquad \forall t \tag{9}$$

$$y_t^g = A_t L_t^{1-\gamma} k_t^{\gamma} \qquad \forall t \tag{10}$$

$$e_t = S_t(1-\mu_t) y_t^g + E_t \qquad \forall t \tag{11}$$

where the vector $\boldsymbol{x} = \{\boldsymbol{o}, \boldsymbol{y}, \boldsymbol{y}^g, \boldsymbol{k}, \boldsymbol{l}, \boldsymbol{e}, \boldsymbol{\tau}, \boldsymbol{u}\}$ in objective (4) is used to refer to all other variables in the formulation, similar to the notation used in the general description in Section 2. Note that elements of $\boldsymbol{x}$ and $\boldsymbol{\mu}$ 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 (4) 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 (5). This constraint defines utility in each period as an isoelastic function of consumption o_t, where β is the calibrated elasticity parameter. Consumption, on the other hand, is defined by equation (6) as the difference between output of goods/services y_t and the capital investment l_t. Capital balance relationship is represented through constraint (7). Constraint (8) represents a set of constraints that link economic activity and the resulting 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 $H(\tau_{t-1}, e_t)$ is increasing in e_t. A key equation in the model is (9), 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 (10). 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 (9), is an increasing function of τ as further discussed in the paragraph below. 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 (11). While abatement has obvious benefits, it is costly as the abatement cost function $c_D(\mu_t)$ in the numerator of (9) is increasing in μ_t. Hence, higher abatement reduces the amount of 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 (6), 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.

The second equation we highlight is (9), which shows how the cost of abatement and the damages from climate change impact economic output y_t in each period t. The cost of abatement in DICE is represented by $c_D(\mu_t) = P_t^{1-\theta} B_t \mu_t^{\theta}$ where $\theta = 2.8$ and $P_t^{1-\theta} B_t$ is a product of two constants with B_t modeling the maximum cost of abatement based on the cost of a "backstop" technology, and P_t representing a possible increase in the costs related to the degree of participation in a given policy. Specifically, this participation factor reflects the fact that in some of the policies considered, not all regions participate in reducing emissions, leading to a higher cost of abatement. A backstop technology in this context is defined as a technology that would serve as a perfect substitute for exhaustible resources. The abatement cost function $c_D(\mu_t)$ will be revised to include the impact of technical change due to R&D investments. Moreover, $D_D(\tau_t) = 1 + \pi\tau_t^2$ in equation (9) 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. These procedures are described in Sections 3.2.3 and 4.1.1 below.

3.2.2. Quantifying the Impact of R&D on Abatement Costs The basic relationships modeled 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 in Figure 1, referred to as '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 after R&D investments on the selected projects. Hence, we need to derive emission abatement cost functions under different success scenarios for implementation in the model. We achieve this through the two phase process described below, which are also discussed in detail by Baker and Solak (2011):

Deriving marginal abatement cost curves. Rather than deriving abatement cost functions directly for each success scenario, we first derive Marginal Abatement Cost Curves (MACs), which reflect the cost of reducing emissions by an additional ton. We then convert these curves to abatement cost functions by integration.

We use 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. 2005). We utilize MACs rather than abatement cost curves as our model input for two reasons. First, the MAC is a key unit of analysis in the decision problem illustrated in Figure 1: 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 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. The MACs were then combined with the elicited success probabilities described in Appendix S3 to get probability distributions over MACs for each possible combination of projects. In other words, each portfolio of technology project investments can be associated with a probability distribution over MACs.

In Figure 2 we show a selection of the MACs that were derived using GCAM. The baseline MAC refers to the case with no R&D-induced technical change as defined by Clarke et al. (2008). We use this baseline MAC to measure the impact of technical change. We demonstrate the impact of three projects (one from each technology, i.e. CCS, nuclear, and solar), if each of them were successful

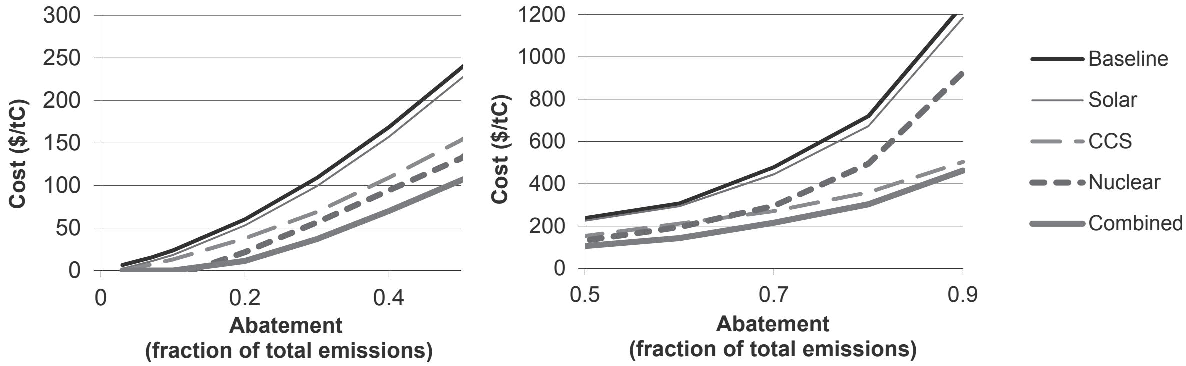

Figure 2 **Representative MACs defining the cost of reducing the carbon emissions by an additional ton. The two plots display the impact of technology projects on the baseline MAC for different ranges of abatement levels.**

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 can be observed to be a combination of pivoting the curve clockwise around zero and shifting the entire curve down. CCS has more of a pivot: 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 generated in GCAM 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 the following equation to estimate a pivot vector $\boldsymbol{\alpha} \equiv (\alpha_{solar}, \alpha_{nuclear}, \alpha_{CCS})$, with each component representing the pivoting effect of the corresponding technology based on the successful projects in that technology, and a shift parameter $h(\boldsymbol{\alpha})$:

$$\widetilde{MAC}(\mu, \boldsymbol{\alpha}) = \prod_i (1-\alpha_i)\, MAC(\mu) - h(\boldsymbol{\alpha})\, MAC(0.5) \tag{12}$$

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.

The process for deriving the values of α_i was as follows. First, a pivot parameter, denoted by α_{ij} was estimated for each individual project j in each technology i. The values of these parameters, which vary depending on the level of success in a project, are listed in Appendix S4. Second, we make the assumption that, within any technology category i, only the best project (the one with the greatest impact on the MAC) will impact the economy. Therefore, we define α_i as $\alpha_i = \max_j \{\alpha_{ij} : j \in \mathcal{S}_i\}$, where the set $\mathcal{S}_i$ refers to any given combination of successful projects j in technology i. Finally, for every combination of possible technological outcomes as represented by the three α_i's for the three technologies, a shift parameter $h(\boldsymbol{\alpha})$ was estimated numerically. Since there is a unique shift value for every combination of α_i values, we represent h as a function of $\boldsymbol{\alpha}$.

The relationship in (12), 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 $\Phi : \mathcal{F} \times \mathbb{R}^n \rightarrow \mathcal{F}$, where $\mathcal{F}$ is a set of functions and $\mathbb{R}^n$ is a set of real vectors, as follows:

$$\Phi(c(\mu), \boldsymbol{\alpha}) = \prod_i (1 - \alpha_i)\, c(\mu) - h(\boldsymbol{\alpha})\, c(0.5)\, \mu \tag{13}$$

This functional can be used to represent 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 below.

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 $\alpha_2 = 1 - (1 - \alpha_{nuc})(1 - \alpha_{sol})$. Thus, we represent R&D investments in the three technologies 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 (6). We assume that R&D investment will take place over the next 50 years at a constant rate. In the base model, output in each period is divided between consumption and investment in traditional capital as noted in equation (6). 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, and thus $t = 5$ implies 50 years:

$$y_t = o_t + l_t + \kappa(\Upsilon_1 + \Upsilon_2)/5 \qquad \forall t \leq 5 \tag{14}$$

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 $\kappa \sum_{i=1}^{2} \Upsilon_i/5$. While Υ_i is a decision variable, the opportunity cost κ is a parameter for which we perform a sensitivity analysis as part of the computational experiments.

Integrating R&D induced technical change into the model: Modification of constraint (9). Following the cost function structure described by equation (13) in Section 3.2.2, we model the impact of technical change by altering constraint (9) as follows:

$$y_t = \frac{1 - \Phi\left(c_D\left(\mu_t\right), \boldsymbol{\alpha}\right)}{D_D\left(\tau_t\right)} y_t^g = \frac{1 - \prod_i (1-\alpha_i)\left[c_D\left(\mu_t\right) - h\left(\boldsymbol{\alpha}\right) c_D\left(0.5\right)\mu\right]}{D_D\left(\tau_t\right)} y_t^g \quad \forall t \tag{15}$$

The model of technical change in equation (15), 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 following quantities:

$$(1-\alpha_1)(1-\alpha_2) \approx 1 - 0.8\alpha_1 - 0.92\alpha_2 \tag{16}$$

$$(1-\alpha_1)(1-\alpha_2)\, h\left(\boldsymbol{\alpha}\right) \approx 0.02 - 0.06\alpha_1 + 0.14\alpha_2 \tag{17}$$

Thus, we express the revised production function in our model for $t > 5$ as follows.

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

We show in Appendix S5 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 depicted in Figure 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 success scenarios impact emission abatement costs. We can calculate the probability of each success scenario from the elicitation data, however, these probabilities are dependent on the technology investment decisions. For example, two different investment portfolios may result in the realization of same pivoting and shifting effects, but the probability of such a realization would be different for each portfolio. This is the endogenous uncertainty in Figure 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 procedures, specifically in stochastic programming. To overcome this issue, we develop a procedure to derive returns functions with fixed probabilities, where each function represents how the values of $\boldsymbol{\alpha}$ change as a function of the investment amount $\boldsymbol{\Upsilon}$.

A reduced-form portfolio model. The procedure we develop to derive returns functions with fixed probabilities is built upon some insights obtained from a reduced form portfolio model that we summarize in Appendix S6. Solving this model leads to two key conclusions reported by Baker and Solak (2011), which are directly related to our analysis in this paper. First, the authors note that the composition of the optimal portfolio 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. This, in turn, results in a probability distribution for the pivot and shift parameters for each investment level. We use these 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

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

Table 1 **Comparison of the estimated and actual probability distribution data over each possible outcome of α_1 for different levels of investment in CCS.**

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 well-rounded and 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, $i = 1, 2$. 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 optimal portfolio of projects for that budget level. Each portfolio is associated with a probability distribution over the possible outcomes of α_i in each technology category. Since there is only one optimal portfolio per budget level, this allows us to associate a probability distribution over the α_i values for each given budget level. We considered a wide range of possible R&D budgets, and chose those that corresponded to 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.

In the lower half of Table 1 we show these values for CCS, with a probability distribution in each column and the α_1 values on the right. However, as noted previously, such endogenous probabilities are typically not amenable to stochastic optimization methods, and thus 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. $\mathcal{A}_i : \Upsilon_i \rightarrow \alpha_i$. The functions $\mathcal{A}_i$ are piecewise linear as they are defined based on a discrete set of investment levels. Tables 2 and 3 show examples of realizations of such returns functions. For example, the bottom row in Table 2 shows a returns function that has a probability of 0.42 of being realized. If this function is realized, then $\alpha_1 = \mathcal{A}_1(\Upsilon_1)$, where

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 2 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.

Budg.($mil)	77	346	423	539	925	1967	3628	4014	4342	8975	20171	Prob.
α_2	0	0	0	0	0	0	0	0	0	0	0	0.087
	0.022	0	0.022	0.022	0.022	0.022	0.022	0.022	0.022	0.022	0.022	0.089
	0.022	0	0.022	0.022	0.022	0.327	0.327	0.327	0.327	0.327	0.327	0.064
	0.022	0	0.022	0.022	0.022	0.131	0.131	0.131	0.131	0.131	0.131	0.044
	0.022	0	0.34	0.34	0.34	0.34	0.34	0.34	0.34	0.34	0.34	0.039
	0.022	0.325	0.325	0.325	0.325	0.325	0.325	0.325	0.325	0.325	0.325	0.157
	0	0.325	0.325	0.325	0.325	0.325	0.327	0.327	0.327	0.327	0.327	0.081
	0	0.325	0.325	0.34	0.34	0.34	0.34	0.34	0.34	0.34	0.34	0.041
	0	0.325	0.34	0.34	0.34	0.34	0.34	0.34	0.34	0.34	0.34	0.124
	0	0.325	0.34	0.34	0.34	0.342	0.342	0.342	0.342	0.342	0.342	0.067
	0.022	0.325	0.325	0.325	0.325	0.325	0.325	0.325	0.325	0.325	0.361	0.01
	0.022	0	0.022	0.022	0.022	0.022	0.022	0.022	0.022	0.134	0.134	0.039
	0	0	0	0	0	0.111	0.111	0.111	0.111	0.342	0.342	0.023
	0	0	0	0	0	0.342	0.342	0.342	0.342	0.342	0.342	0.023
	0	0	0	0	0	0.327	0.342	0.342	0.342	0.342	0.342	0.025
	0	0	0	0	0	0	0.327	0.361	0.361	0.361	0.361	0.028
	0	0	0	0	0	0	0	0	0.342	0.115	0.115	0.037
	0	0	0	0	0	0.325	0.325	0.022	0.34	0.34	0.34	0.022

Table 3 Piecewise linear returns functions for solar-nuclear, where the central columns show values of α_2 for discrete levels of investment. Each row, which corresponds to a realization of the function $\mathcal{A}_2$, is associated with a probability given in the far right column.

$$\mathcal{A}_1(\Upsilon_1) = \begin{cases} 0 & \text{if } \Upsilon_1 < 52 \\ 0.319 & \text{if } 52 \le \Upsilon_1 < 108 \\ 0.38 & \text{if } \Upsilon_1 \ge 108 \end{cases} \tag{19}$$

The set of α_i values shown in Tables 2 and 3 correspond to the possible set of realizations for each combination of projects, and are derived from the parameters α_{ij} for each project j, which are listed in Appendix S4.

While it is possible to enumerate all possible returns functions, this implies a very large number of realizations, which is not tractable.. Thus, we perform a scenario reduction process and identify a subset of the possible returns functions with a good approximation of the actual distribution. This process is 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.

Tables 2 and 3 show the estimated returns functions for the CCS and solar-nuclear technology categories, respectively. For example, in Table 3 the fifth row from the bottom represents a returns function that has $\alpha_2 = \mathcal{A}_2(\Upsilon_2) = 0$ if $\Upsilon_2 < \$1,967$ million, and $\alpha_2 = \mathcal{A}_2(\Upsilon_2) = 0.342$ if $\Upsilon_2 \geq \$1,967$ million. The probability that this particular function is realized is 0.023. 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 1 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% in each case. Similar results hold for the approximations of the solar-nuclear returns functions as well, with differences being less than 4% in each case.

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 a set of new variables and constraints. The details of this implementation are described in Appendix S6.

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, 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 Tables 2 and 3, and a π value of 1. Assuming that the latter can occur with probability of 1 and using the probabilities shown in Tables 2 and 3, 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)-(3) 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 $\boldsymbol{\Upsilon}, \boldsymbol{\mu}_N, \boldsymbol{x}_N$ by possibly different vectors $\boldsymbol{\Upsilon}^\omega, \boldsymbol{\mu}_N^\omega, \boldsymbol{x}_N^\omega$, similar to the 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. 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. 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. 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 (4)-(11). However, we show in Appendix S5 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 the solution methodology for the problem.

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

$$\max_{\boldsymbol{\Upsilon}^{\Omega},\boldsymbol{\mu}^{\Omega},\boldsymbol{x}^{\Omega}} \sum_{\omega\in\Omega} p^{\omega} \sum_{t} R_t u_t^{\omega} \tag{20}$$

$$\text{s.t. } J_{\psi}^{\omega}(\boldsymbol{\Upsilon}^{\omega},\boldsymbol{\mu}^{\omega},\boldsymbol{x}^{\omega}) \le b_{\psi}^{\omega} \qquad \forall \psi,\omega \tag{21}$$

$$\Upsilon_i^{\omega} - \sum_{\omega'\in\Omega} p^{\omega'} \Upsilon_i^{\omega'} = 0 \qquad \forall i,\omega \tag{22}$$

$$k_t^{\omega} - \sum_{\omega'\in\Omega} p^{\omega'} k_t^{\omega'} = 0 \qquad \forall t \le 5,\omega \tag{23}$$

$$u_t^{\omega} - \sum_{\omega'\in\Omega} p^{\omega'} u_t^{\omega'} = 0 \qquad \forall t \le 5,\omega \tag{24}$$

where the superscripts in $\boldsymbol{\Upsilon}^{\Omega}$, $\boldsymbol{\mu}^{\Omega}$, and $\boldsymbol{x}^{\Omega}$ denote that the optimization is performed over all scenarios in the scenario set Ω. In this formulation, objective (20) corresponds to the objective function (1) in the general formulation and represents the maximization of expected total social utility over all scenarios. Constraints (21) 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 (5), (7), (8), (10), (11), as well as the extended relationships modeled by equations (14), (18) and the required constraints for the piecewise linear mappings described in Appendix S7. 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 (22)-(24) 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 (21)-(24) together define a reformulation of constraints (2)-(3) in the general formulation described in Section 2.

Model (20)-(24) is a two-stage stochastic nonlinear programming problem that can be solved through decomposition methods, as the size of the scenario set does not allow for the direct solution of the problem. As our 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 S8.

5. Experimental Setup for R&D Policy Analysis

In this section we discuss the policy experiments we ran with our stochastic optimization model. 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 may not, in fact, be the actual cost to society. Additional costs to society include administrative costs to get the money into the hands of scientists and engineers, inefficiencies in the funding allocation process, and the fact that money spent on one particular type of R&D tends to reduce other directions of R&D. 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 (14) 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 as part of 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. They include both low and high abatement levels, different levels of flexibility, and different assumptions about regional participation. Here we describe these policies, which are also summarized in Table 4. For each 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 and growth 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

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.

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[3].

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 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 simply adds a single constraint that limits the average global temperature increase to 2°C. In our modeling framework this constraint is only minimally binding.

We note here that we are not evaluating the different policy environments against one another. We recognize that the DICE Optimal policy will, by definition, be optimal within the framework

[3] 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.

Investments ($ million)									Total Inv ($bil)
CCS			Nuclear			Solar			
Pre C	Chem L	Post C	LWR	HTR	FR	Org	Inorg	3rd G	
386	56	519	346	3089	15443	830	77	386	21.132
386	56	519	346	3089	0	830	77	0	5.303
154	56	519	346	3089	0	830	77	0	5.071
154	56	519	346	1544	0	0	77	0	2.696

Table 6 Allocation of total investment under different optimal investment values.

considered. Nor do we intend to make a judgement about the appropriate discount rate. Rather, this analysis is intended to evaluate the role of R&D within different policy environments.

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 (9). 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). For this paper we have chosen to work with 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.

6.1. Optimal Investment in R&D

We find that *the optimal investment in the R&D projects considered in this analysis is quite robust, both across different policy environments and across different risk cases.* Figure 3 illustrates the optimal investment across risk and policy environments. The horizontal axis represents risk cases $1-3$ from Table 5. The value on the horizontal axis is the GDP loss given a 2 degree increase in temperature for

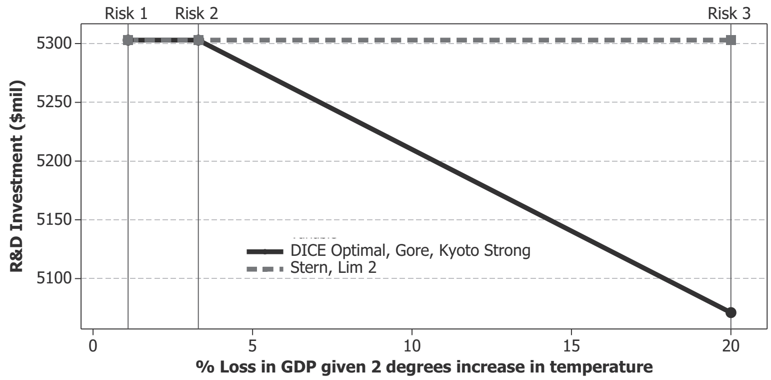

Figure 3 **Optimal investment across risk and policy scenarios, where the horizontal axis values correspond to the GDP loss for the high damage outcome of risk cases 1 – 3 from Table 5.**

the high damage outcome for each risk case, considering the no-, medium-, and high-risk cases. Thus, as we move to the right, damages are getting more uncertain. 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. 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 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. We discuss further why the overall investment is so robust in Section 6.3 below.

In Table 6, we show the allocation of the total investment in research areas for the optimal levels of investment in different cases. For example, an investment of $5,303 million represents a high investment in all CCS projects, as well as the light water reactor (LWR), high temperature reactor (HTR), organic and inorganic solar projects. This is the optimal investment for Stern and Lim 2 at all risk levels, plus for DICE Optimal, Gore, and Kyoto Strong for risk cases 1 and 2. For DICE Optimal and the other two, the optimal investment reduces slightly under high risk (case 3), to $5,071. The intermediate risk case also has this investment level for DICE Optimal.

Not shown in Figure 3 is the DICE Optimal investment under very high risk. This falls to $2,696 million, representing high investments in chemical looping, post-combustion, LWR, and inorganic solar, as well as medium investment in pre-combustion and HTR. In general, we see a somewhat monotonic response to risk in these 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

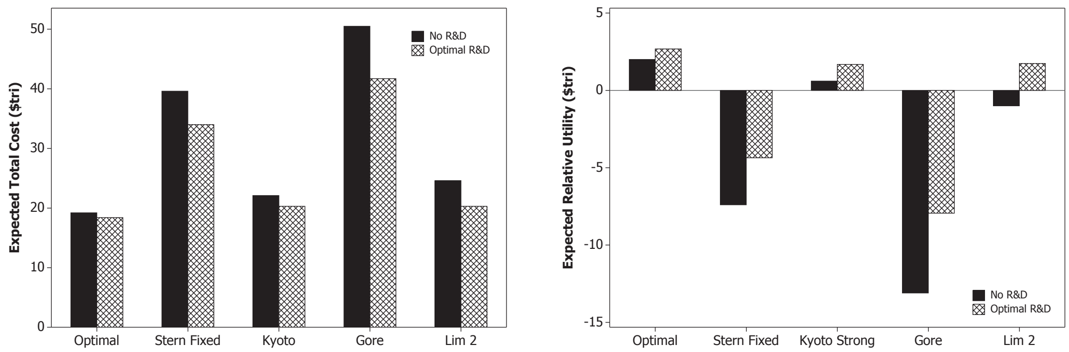

(a) Effect of R&D on total costs (b) Expected relative utility of policy interventions

Figure 4 Expected costs and utilities of policy interventions

100% with or without technical change: increasing damages further than this has no impact on the total cost of abatement. 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. We find that 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.

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 4, we illustrate the impact of R&D on the different policy environments. Figure 4(a) shows the expected total cost of each policy with and without R&D, while Figure 4(b) shows the expected relative utility of each policy intervention with respect to Baseline with and without R&D.

The vertical axis in Figure 4(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 4(b) represents expected utility translated into trillions of 2005 dollars. Each bar represents the extra utility gained (or lost) from the Baseline by implementing the policy intervention. The darker bars (the left bar for each policy) are in the absence of R&D and are very similar to the results in Nordhaus (2008). The lighter bars (the right bar for each policy) are when R&D is available and chosen optimally. Note that some of the policy environments are an improvement over doing nothing, whereas some are worse than doing nothing, at least as evaluated within DICE. 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 absolute 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 representation of a possibly implementable policy. In the absence of R&D, it is barely better than doing nothing when evaluated in the DICE model. 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 5(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. The bars for each policy are displayed in the same order in the figure. 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 3. The higher and lower numbers increase and decrease these central values by $386 million and $328 million, respectively. Figure 5(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*. Note that, for a comparative illustration, the utilities of the three risk cases in Figure 5(b) have been normalized so that they are equal at the $5,071 portfolio level.

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

Figure 6(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. For example, the probability of having optimal abatement between 68–77% in 2105 is 45%. This is associated with scenarios in which there is success in both nuclear and CCS. Then next cluster, with probability of 52%, includes scenarios in which either nuclear or CCS fails. The lowest line is the case where all technology fails. In Figures 6(b) and 7 we will just show the range of paths without associated probabilities, since they follow a similar pattern.

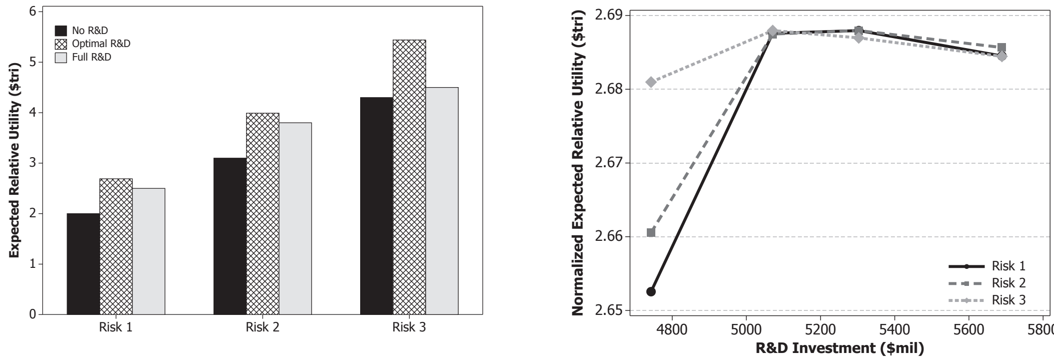

(a) Expected utility of policy intervention for the DICE Optimal policy for no and full R&D budgets

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

Figure 5 Expected utility of policy intervention.

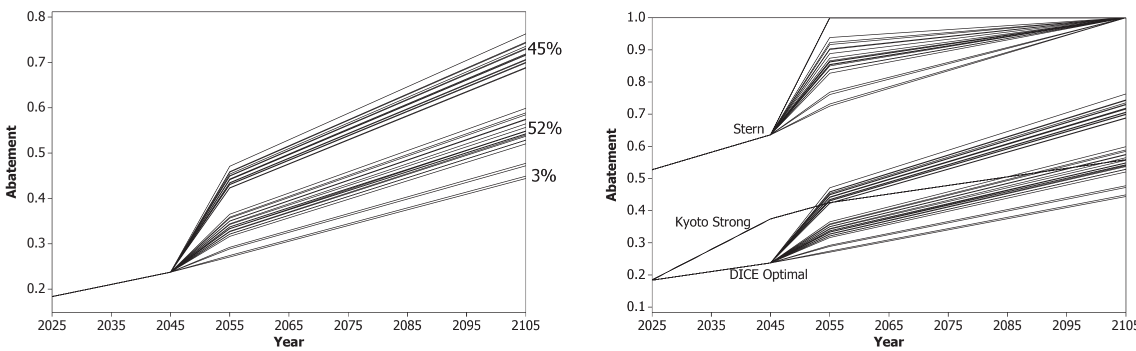

(a) Range of abatement paths in the risk 1 case for DICE Optimal

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

Figure 6 Ranges of abatement paths for different realizations of technical uncertainties.

Figure 6(b) compares the range of emissions paths from the DICE Optimal, Stern, and Kyoto Strong policies. The high group are the paths from Stern, the lower group from DICE Optimal, and the middle line is for Kyoto Strong.

Note that in the absence of technical change the Kyoto Strong path is clearly 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.*

Figure 7(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 (the higher group) and for Stern (the lower group). From this figure we can make three observations. First, all the DICE Optimal paths are above 2 degrees between 2075 and 2200, while all Stern paths are always below. Note that the DICE model appears to be relatively optimistic about the ability to stay below 2 degrees. A number of large IAMs

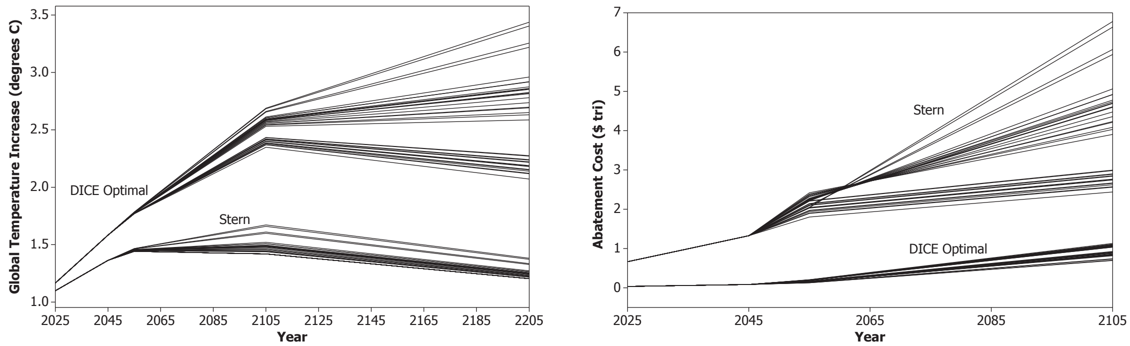

(a) Range of temperature paths in the risk 1 case for DICE Optimal and for Stern

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

Figure 7 Ranges of temperature and abatement cost paths.

find this goal to be infeasible to solve (Clarke et al. 2009), indicating it is very difficult. Thus, these results must be interpreted within the DICE framework. 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. R&D can only impact the temperature to the degree that it impacts abatement. Since abatement is already very high in Stern, R&D has only a limited impact. In DICE Optimal, on the other hand, abatement in 2105 ranges from 44% to 74% depending on the outcome of R&D.

Finally, 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 7(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 with more reduction in costs. In fact, Figure 7(b) is almost the opposite of Figure 7(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.

Figure 8 illustrates this idea, showing MAC curves with and without technical change, i.e. MAC and $\widetilde{MAC}$, respectively. The horizontal line represents marginal damages. At the abatement level represented by the point a, the abatement cost before technical change is the area oad. If abatement switches to b after technical change, the new abatement cost is the area obc. The cost-side benefit of

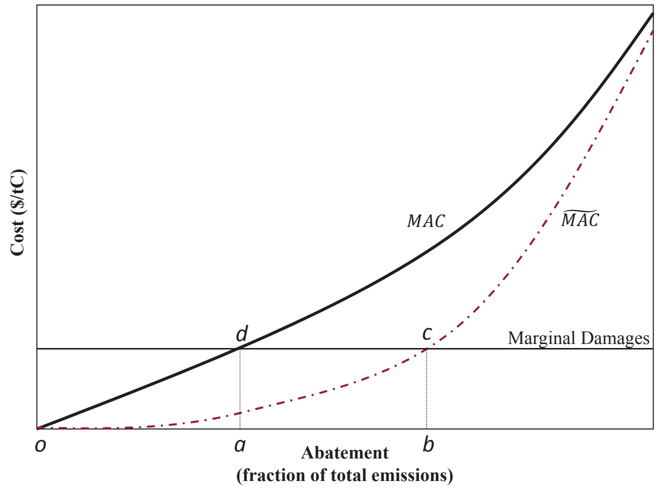

Figure 8 **Impact of R&D on MAC curves.**

technical change is $oad - obc$. Note, however, that this “benefit” can be negative – abatement costs may be higher after technical change than before. The environmental-side benefit of moving abatement from a to b is the area of the rectangle $abcd$. If abatement is relatively low and adjusts optimally (the case illustrated in this figure), then it can be seen that obc may in fact be larger than oad – there is no cost-side benefit. But, because abatement adjusts considerably, there is a considerable environmental side benefit. As the pre-technical change abatement increases and the point a moves to the right, cost side benefits of technical change get larger and environmental side benefits of technical change get smaller. This is similar to what we see in the Stern policy, and also in cases of high damages.

The robustness of the 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 ends up having 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. Thus, the optimal investment is surprisingly robust to the different environments.

6.4. R&D and Riskiness of Outcomes

In this section we look at how investment into R&D impacts the riskiness of the policy outcomes. Figure 9 shows part of three cumulative distribution functions (CDFs). The CDFs are for DICE Optimal under high risk (risk 3), comparing no R&D, optimal R&D, and full R&D. The horizontal axis represents the present value of total costs (cost of abatement plus cost of damages). Each point on the graph represents the probability that total costs are less than or equal to the value on the horizontal axis. For example, the probability that the total cost is less than $170 trillion, given an optimal investment in R&D, is about 0.98. We only show the far right of the graph, since there is no visual difference between the three cases on the rest of the graph. Note that society would prefer to

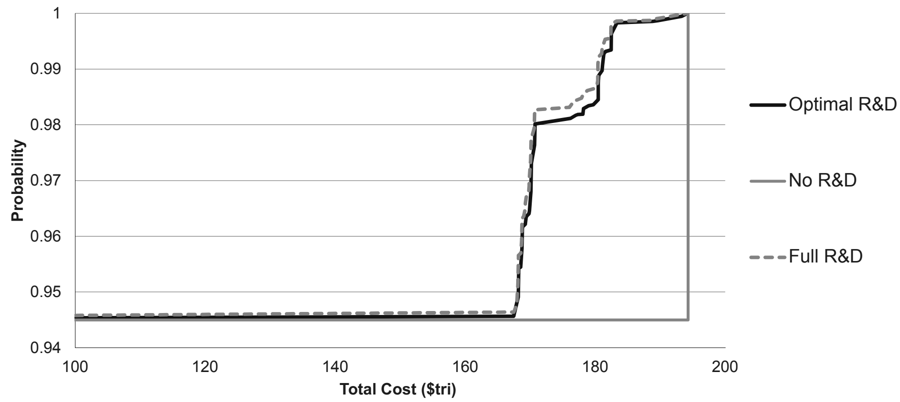

Figure 9 Cumulative distribution function comparing different R&D policies in the risk 3 DICE Optimal case.

be as far left as possible on this graph, and so a higher line is preferred to a lower line. There is a 5.5% chance that high damages (about 20 times higher than the mean) realize in the risk 3 case. If there is no R&D, then damages in this case will be equal to $194 trillion. With full or optimal R&D however, damages may be limited, with only about a 2% chance that damages are greater than $170 trillion. Thus, *R&D provides risk reduction* (visualized as the area between the darker and the lighter solid curves).

There is very little difference between the CDFs for optimal R&D and full R&D. Full R&D provides some risk reduction in the high risk case, and thus lies above the optimal R&D line at some places.

7. Conclusions and Further Policy Insights

Finding an optimal R&D portfolio in the face of climate change is a very challenging problem, and the data that are available is sparse and not as airtight as we would like. It is, however, a real and pressing problem faced by the U.S. (in particular the U.S. Department of Energy) 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 detailed and dynamic stochastic model. Thus, this work provides a framework which can be updated as better information becomes available, such as more detailed elicitations about these technologies, data on other technologies, or updated model results.

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 policy and the DICE Optimal policy are strikingly 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, we cannot guarantee robust results for any set of data on any set of technologies. 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 affects, 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 (i.e. that the probability distribution is not known), these 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. 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 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.

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 (such as technological change) 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 convex representation of a highly nonlinear stochastic problem. Several characteristics of this complex problem were exploited enabling a fast and effective solution procedure. 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).

References

Anadon, L., V. Bosetti, M. Bunn, M. Catenacci, A. Lee. 2011a. Report of international workshop on RD&D to enhance the role of nuclear energy in meeting climate and energy challenges. Tech. rep., Harvard Kennedy School.

Anadon, L., M. Bunn, G. Chan, M. Chan, C. Jones, R. Kempener, A. Lee, N. Logar, V. Narayanamurti. 2011b. Transforming U.S. energy innovation: Report for energy technology innovation policy research group. Tech. rep., Harvard Kennedy School.

Baker, E. 2009. Uncertainty and learning in climate change. *J. of Public Economic Theory* **11** 721 – 747.

Baker, E., H. Chon, J. Keisler. 2008. Nuclear Power: Combining expert elicitations with economic analysis to inform climate policy. Available at SSRN: http://ssrn.com/abstract=1407048.

Baker, E., H. Chon, J. Keisler. 2009a. Advanced solar R&D: Combining economic analysis with expert elicitations to inform climate policy. *Energy Economics* **31** S37 – S49.

Baker, E., H. Chon, J. Keisler. 2009b. Carbon capture and storage: Combining expert elicitations with economic analysis to inform climate policy. *Climatic Change* **96(3)** 379.

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**.

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

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

Blanford, G.J., J.P. Weyant. 2007. Optimal investment portfolios for basic R&D. Working Paper, Stanford University.

Bosetti, V., M. Catenacci, G. Fiorese, E. Verdolini. 2012. The future prospect of PV and CSP solar technologies: An expert elicitation survey. Working papers, Fondazione Eni Enrico Mattei.

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** 37–45.

Chan, G., L. Anadon, M. Chan, A. Lee. 2011. Expert elicitation of cost, performance, and RD&D budgets for coal power with CCS. *Energy Procedia* **4** 2685–2692.

Clarke, L., J. Edmonds, V. Krey, R. Richels, S. Rose, M. Tavoni. 2009. International climate policy architectures: Overview of the EMF 22 international scenarios. *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: An updated analysis of advanced technology scenarios. Tech. Rep. PNNL-18075, Pacific Northwest National Laboratory.

Curtright, A., M. Morgan, D. Keith. 2008. Expert assessments of future photovoltaic technologies. *Environmental Science & Technology* **42** 9031–9038.

Edmonds, J.A., J.F. Clarke, J.J. Dooley, S.H. Kim, S.J. Smith. 2005. Stabilization of CO_2 in a B2 world: Insights on the roles of carbon capture and storage, hydrogen, and transportation technologies. John Weyant, Richard Tol, eds., *Special Issue, Energy Economics*.

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

Hiriart-Urruty, J-B., C. Lemarechal. 1993. *Convex Analysis and Minimization Algorithms I*. Springer.

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

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. http://www.nap.edu/catalog/11806.html.

Nordhaus, W.D. 1993. Rolling the DICE: An optimal transition path for controlling greenhouse gases. *Resource and Energy Economics* **15** 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 and IIASA, Washington D.C. and Laxenburg, Austria, 182–209.

Nordhaus, W.D. 2008. *A Question of Balance: Economic Modeling of Global Warming*. Yale University Press, New Haven, CT.

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

Pizer, W. A., D. Popp. 2008. Endogenizing technological change: Matching empirical evidence to modeling needs. *Energy Economics* **30** 2754–2770.

Rao, A., E. Rubin, D. Keith, M. Morgan. 2006. Evaluation of potential cost reductions from improved amine-based CO_2 capture systems. *Energy Policy* **34** 3765 – 3772.

Solak, S., E. Baker. 2012. Convexity analysis for the DICE model. Manuscript under Review.

Stern, N. 2007. *The Economics of Climate Change: The Stern Review*. Cambridge University Press, Cambridge, MA.

Tversky, A., D. Kahneman. 1974. Judgement under uncertainty: Heuristics and biases. *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. Pivot Parameter Values for Individual Technology Projects

Appendix S5. Proof of Analytical Results

Appendix S6. Reduced Form R&D Model

Appendix S7. Representation of Stochastic Returns Functions

Appendix S8. Description of the Solution Procedure

Appendix S1. Definitions of Acronyms

$3rdG$: 3rd Generation Technologies
CCS: Carbon Capture and Storage
CDF: Cumulative Distribution Function
$ChemL$: Chemical Looping
$DICE$: Dynamic Integrated Model of Climate and the Economy
FR: Fast Burner Reactors
$GCAM$: Global Change Assessment Model
GDP: Gross Domestic Product
HTR: High Temperature Reactors
IAM: Integrated Assessment Model
$Inorg$: Inorganic Solar Cells
LWR: Light Water Reactors
MAC: Marginal Abatement Curve
Org: Organic Solar Cells
$PostC$: Post-combustion Carbon Capture
$PreC$: Pre-Combustion Carbon Capture

Appendix S2. Definitions of Variables and Parameters

Variables

$c(\mu_t)$: generic representation for cost of emissions abatement
$c_R(\mu)$: cost of emissions abatement in the reduced form R&D model
$D_R(\mu)$: cost of climate damages in the reduced form R&D model
$c_D(\mu)$: cost of emissions abatement in the DICE model
$D_D(\tau_t)$: cost of climate damages in the DICE model
e_t : total carbon emissions in period t
$\boldsymbol{G}_s(\cdot)$: representative function modeling the constraint set for stage s, $s = N, L$
$H(\tau_{t-1}, e_t)$: representative function linking economic activity and greenhouse gas emissions to the carbon cycle
h: shift effect in the marginal abatement curve due to technological success
$J_\psi^\omega(\boldsymbol{\Upsilon}^\omega, \boldsymbol{\mu}^\omega, \boldsymbol{x}^\omega)$: representative function modeling the constraint ψ for scenario ω
k_t : capital stock in period t
l_t : investment in traditional capital in period t
o_t : consumption of goods/services in period t
u_t : social utility in period t
$U_s(\cdot)$: social utility in stage s, $s = N, L$
x_{ijk}: 1 if project j of technology i is funded at level k in the reduced form model, 0 otherwise
$\boldsymbol{x}_s$: generic vector representing all decision variables other than abatement decisions in stage s, $s = N, L$
$\boldsymbol{x}$: generic vector representing all decision variables other than abatement decisions
y_t : net output of goods/services in period t
y_t^g : unadjusted output in period t
α_i: pivot effect in MAC due to success in technology i
μ_t : level of emissions abatement in period t
$\boldsymbol{\mu}_s$: vector representing emissions abatement decisions in stage s, $s = N, L$
$\Phi\left(c\left(\mu\right), \boldsymbol{\alpha}\right)$: cost of emissions abatement after technical change
τ_t : atmospheric temperature in period t
Υ_i: investment in technology category i

Υ: vector representing investment decisions in technologies

Parameters

$\mathcal{A}_i$: stochastic parameter representing the returns function for technology category i
A_t : level of total factor productivity in period t
b_ψ^ω: representative parameter modeling the right hand side of constraint ψ for scenario ω
$\boldsymbol{b}_s$: representative vector modeling the the right hand sides of constraints for stage s, $s = N, L$
$\mathcal{B}$: R&D budget
B_t : maximum cost of abatement based on the cost of a backstop technology in period t
E_t : emissions from deforestation in period t
f_{ijk}: required investment for project j of technology category i at level k in the reduced form R&D model
L_t : population and labor input in period t
p^ω: probability of scenario ω
P_t : degree of policy participation in period t
R_t : utility discount factor for period t
S_t : ratio of uncontrolled emissions to output in period t
Z: stochastic parameter modeling the magnitude of climate damages in the reduced form R&D model
α_{ij}: calculated pivot effect in MAC due to success in project j in technology i
β : elasticity of marginal utility of consumption
γ : elasticity of output with respect to capital
θ : cost function exponent set to 2.8 in DICE
κ : opportunity cost parameter
π : stochastic parameter modeling the magnitude of climate damages
σ : rate of depreciation of capital

Appendix S3. Summary of Expert Elicitation Results

Technology	Project	NPV of Funding (000,000)	Probability of success
Carbon Capture and Storage (CCS)	*Pre-combustion Carbon Capture (Pre C)*	$39	2.7%
		$154	11.0%
		$386	22.3%
	Chemical Looping (Chem L)	$19	8.0%
		$38	29.5%
		$56	42.0%
	Post-combustion Carbon Capture (Post C)	$52	59.0%
		$224	70.0%
		$519	78.5%

Table S3.1 Summary of expert elicitation results for the CCS technology (Baker et al. 2009b).

Technology	Project	NPV of Funding (000,000)	Probability of success
Nuclear	*Light Water Reactors (LWR)*	$173	21.3%
		$260	33.8%
		$346	60.0%
	High Temperature Reactors (HTR)	$772	0.3%
			1.2%
		$1,544	17.0%
			26.2%
		$3,089	30.2%
			40.3%
	Fast Burner Reactors (FR)	$1,158	0.1%
			7.5%
		$4,633	0.5%
			32.5%
		$15,443	16.3%
			60.0%

Table S3.2 Summary of expert elicitation results for the nuclear technology (Baker et al. 2008).

Technology	Project	NPV of Funding (000,000)	Probability of success
Solar	*Organic Solar Cells (Org)*	$116	0.0%
			13.0%
		$830	3.9%
			28.7%
	Inorganic Solar Cells (Inorg)	$39	26.7%
		$77	44.3%
	3rd Generation Technologies (3rd G)	$386	2.0%

Table S3.3 Summary of expert elicitation results for the solar technology (Baker et al. 2009a).

Different research areas in each technology are listed in the 'Project' column in the elicitation summary tables. The investment amount for each project can be at one of multiple potential levels. These different investment levels are listed under the 'NPV of Funding' column in the tables, where NPV refers to the net present value calculated at an interest rate of 5%. Each project is also associated with specific endpoints or targets to be assessed, such as a given efficiency level, which define "success" for that project. The specific probabilities of success for different investment levels of each project reflect an aggregation of the individual experts' judgments, and are shown in the last column of the tables. Note that for some of the technologies, two levels of success were defined (representing lower and higher goals such as 15% versus 31% efficiency), and therefore the funding amounts have two rows associated with them.

Appendix S4. Pivot Parameter Values for Individual Technology Projects

Technology	Project	Pivot Parameter
Solar	*Organic Solar Cells*	0.050
		0.022
	Inorganic Solar Cells	0.022
	3rd Generation Technologies	0.050
Carbon Capture and Storage	*Pre-comb. Carbon Capture*	0.347
	Chemical Looping	0.380
	Post-comb. Carbon Capture	0.319
Nuclear	*Light Water Reactors*	0.325
	High Temperature Reactors	0.327
		0.111
	Fast Burner Reactors	0.332
		0.115

Table S4.1 Pivot parameter values for individual technology projects. Multiple parameter values for a project imply that values may differ based on the funding level.

Note that the combined solar/nuclear parameter α_2 is calculated as $\alpha_2 = 1-(1-\alpha_{nuc})(1-\alpha_{sol})$. Also note that for each combination of α_1 and α_2 there is a unique $h(\alpha_1, \alpha_2)$. For these values please contact the authors.

Appendix S5. Proof of Analytical Results

Lemma 1 **(Convexity of equation** (18)**).** *The revised output equation* (18) *can be expressed as a convex inequality constraint.*

Proof: Solak and Baker (2012) show that equation (9) in DICE can be expressed as

$$y_t - \frac{1-c_D(\mu_t)}{D_D(\tau_t)} y_t^g \le 0 \tag{25}$$

and that the left hand side of this constraint is convex in the decision variables for the range of parameter values used in DICE. This implies that the function is concave in the numerator $1-c_D(\mu_t)$. Let scalar function $\ell : \mathbb{R} \to \mathbb{R}$ be defined such that

$$\ell(r) = -\frac{r}{D_D(\tau_t)} y_t^g \tag{26}$$

Hence, equation (18) can be expressed as $y_t + \ell(\varrho(\mu_t, \alpha_1, \alpha_2)) \le 0$, where

$$\varrho(\mu_t, \alpha_1, \alpha_2) = [1-((1-0.8\alpha_1-0.92\alpha_2)c_D(\mu_t)-(0.02-0.06\alpha_1+0.14\alpha_2)c_D(0.5)\mu_t)] \tag{27}$$

Hence for the convexity of (18), it suffices to show that $\ell(\varrho(\mu_t, \alpha_1, \alpha_2))$, i.e. the composition of ℓ and ϱ is convex in the decision variables $\mu_t, \alpha_1, \alpha_2$.

Note that composition of a function with a scalar convex function is convex if the function is concave and the extended-value extension of the scalar function is nonincreasing. Given that ℓ is nonincreasing, we need to show that $\varrho(\mu_t, \alpha_1, \alpha_2)$ is concave. This can be shown by computing the Hessian of the function, and noting that the Hessian is negative semidefinite, which we skip the details for. It follows that (18) has an equivalent convex representation. □

THEOREM 1 **(Convexity of the R&D portfolio optimization model)**. *The stochastic programming formulation* (20)-(24) *for the stochastic R&D portfolio optimization model is convex with respect to all decision variables.*

Proof: The result follows from the proof of convexity for the deterministic DICE model by Solak and Baker (2012), Lemma 1, and the linearity of constraints (14), (22)-(24), and (38)-(40). □

PROPOSITION 1 **(Sufficiency of nonanticipativity in Υ, k, and u)**. *Let Υ_i^ω, k_t^ω, u_t^ω, and $\mathbf{x}_t^\omega$ represent the optimal decision variable values for scenarios $\omega \in \Omega$ in the stochastic R&D portfolio optimization model, where $\mathbf{x}_t^\omega$ is the vector of all other variables. For any $\omega, \omega' \in \Omega$, if $\Upsilon_i^\omega = \Upsilon_i^{\omega'}$, $k_t^\omega = k_t^{\omega'}$, and $u_t^\omega = u_t^{\omega'}$, then there exists an optimal solution where $\mathbf{x}_t^\omega = \mathbf{x}_t^{\omega'}$.*

Proof: The result can be established by analyzing the implied relationships in formulation (4)-(8). We first note that the representative constraints (8) involve the following three constraints:

$$m_t^a = e_t + 0.811 m_{t-1}^a + 0.097 m_{t-1}^u \qquad \forall t \tag{28}$$

$$f_t = 3.8 \log\{m_t^a + m_{t+1}^a / 1192.8\} \qquad \forall t \tag{29}$$

$$\tau_t = \tau_{t-1} + 0.22(f_t - 1.27\tau_{t-1} - 0.3(\tau_{t-1} - \underline{\tau}_{t-1})) \qquad \forall t \tag{30}$$

$$\underline{\tau}_t = \underline{\tau}_{t-1} + 0.05(\tau_{t-1} - \underline{\tau}_{t-1}) \qquad \forall t \tag{31}$$

where m_t^a and m_t^u are the carbon concentrations in the atmosphere and upper oceans, f_t is the total radiative forcing, and $\underline{\tau}_t$ is the ocean temperature in period t. The conditions $\Upsilon_i^\omega = \Upsilon_i^{\omega'}$, $k_t^\omega = k_t^{\omega'}$, and $u_t^\omega = u_t^{\omega'}$ have the following implications. First, given the equality $u_t^\omega = u_t^{\omega'}$ and the definition of u_t in constraint (5), we get $o_t^\omega = o_t^{\omega'}$. Similarly, $k_t^\omega = k_t^{\omega'}$ implies through constraint (7) that $l_t^\omega = l_t^{\omega'}$. Moreover, due to the equality of variables o_t, l_t, and Υ_i in constraint (14) for scenarios ω and ω', we note that the condition $y_t^\omega = y_t^{\omega'}$ must also hold. The last relationship, along with the condition $k_t^\omega = k_t^{\omega'}$ requires that the following must hold for scenarios ω and ω':

$$\frac{1 - P_t^{1-\theta}(\mu_t^{\omega'})^\theta B_t}{1 - P_t^{1-\theta}(\mu_t^{\omega})^\theta B_t} = \frac{1 + \pi(\tau_t^{\omega'})^2}{1 + \pi(\tau_t^{\omega})^2} \tag{32}$$

Clearly, this condition will be satisfied when $\mu_t^\omega = \mu_t^{\omega'}$ and $\tau_t^\omega = \tau_t^{\omega'}$, implying the equalities $e_t^\omega = e_t^{\omega'}$ due to constraint (11), and $\underline{\tau}_t^\omega = \underline{\tau}_t^{\omega'}$ due to constraint (31). Based on this, the relationship in (30) requires that $f_t^\omega = f_t^{\omega'}$, and in turn $m_t^{a,\omega} = m_t^{a,\omega'}$ due to constraint (29). Finally, equality of values for variables e_t and m_t^a in constraint (28) results in the condition $m_t^{u,\omega} = m_t^{u,\omega'}$. Hence, it follows that there exists an optimal solution where all variables that are not explicitly included in the nonanticipativity constraints are also equal for ω and ω'. □

Appendix S6: Reduced Form R&D Model

For the reduced form R&D model, we use the simplistic model of Baker and Solak (2011), where the authors reduce the economy into two periods and a single equation. A general representation of this model is as follows:

$$\min_{x_{ijk}:\forall i,j,k} \quad E_{\boldsymbol{\alpha},Z}\left[\min_{\mu}\left[\Phi\left(c_R\left(\mu\right),\boldsymbol{\alpha}\right) + ZD_R\left(\mu\right)\right]\right] \tag{33}$$

$$\text{s.t.} \quad \sum_i \sum_j \sum_k f_{ijk} x_{ijk} \leq \mathcal{B} \tag{34}$$

$$\sum_k x_{ijk} \leq 1 \qquad \forall i,j \tag{35}$$

$$0 \leq \mu \leq 1 \tag{36}$$

$$x_{ijk} \in \{0,1\} \qquad \forall i,j,k \tag{37}$$

The reduced form model determines the abatement level μ and binary technology selection decisions that minimize the expectation of the sum of abatement costs $\Phi\left(c_R\left(\mu\right),\boldsymbol{\alpha}\right)$ and damage costs $D_R(\mu)$. In this objective function representation, $c_R\left(\mu\right) = b_0\mu^{b_1}$ denotes the baseline abatement cost function used in the reduced-form model, where b_0 and b_1 are calibrated parameters. The damage cost function in the reduced-form model is defined as $D_R(\mu) = M_0(Q - M_1\mu)^2$, where Q, M_0, and M_1 correspond to specific parameter values. The technology selection decisions are made in the first period and abatement is performed in the second period after realization of the uncertain parameters, which consist of the technical change indicators $\boldsymbol{\alpha}$ and the magnitude of climate change damages, i.e. Z. The probability distributions over the parameters $\boldsymbol{\alpha}$ and thus $h\left(\boldsymbol{\alpha}\right)$ depend on the R&D projects that are chosen, while the uncertainty around the magnitude Z of climate change damages is exogenous. The indices i, j and k represent the technology (CCS, nuclear, and solar), the specific project for a technology and the level of investment, respectively. The binary decision variables x_{ijk} equal 0 if there is no investment at funding level k in project j of technology i, and 1 otherwise. The other decision variable is abatement $\mu \in [0,1]$, i.e. the fraction of emissions reduced below a business-as-usual level. The investment decisions are constrained by the R&D budget $\mathcal{B}$, and by the fact that a project can be invested in only at one level, where f_{ijk} is the NPV of funding level k for project j of technology i, as given in the third column of the tables in Appendix S3.

Appendix S7. Representation of Stochastic Returns Functions

The stochastic returns functions in the R&D optimization model are represented through a piecewise linear structure. To this end, we define new variables $\lambda_i^n \geq 0$ for $i = 1, 2$ and $n = 0, \ldots, N_i$, where N_i is the number of vertices used to represent the returns functions for technology category i, and then include the following constraints in our formulation:

$$\Upsilon_i^\omega = \sum_{n=1}^{N_i} v_i^n \lambda_i^{\omega n} \qquad \forall i, \omega \tag{38}$$

$$\alpha_i^\omega = \sum_{n=1}^{N_i} \hat{\alpha}_i^\omega(n) \lambda_i^{\omega n} \qquad \forall i, \omega \tag{39}$$

$$\sum_{n=0}^{N_i} \lambda_i^{\omega n} = 1 \qquad \forall i, \omega \tag{40}$$

where v_i^n is the investment parameter value for the nth vertex. The stochastic parameter $\hat{\alpha}_i^\omega(n)$ in these constraints corresponds to the value of the return parameter α_i^ω at the nth vertex of the return function. Note that we must require that at most two adjacent $\lambda_i^{\omega n}$ can be nonzero for each i and ω to ensure that corresponding values of Υ_i^ω and α_i^ω lie on one of the straight line segments of the returns function. However, this condition is satisfied regardless due to the result in Appendix S7 that our stochastic R&D optimization model is convex.

Appendix S8. Description of the Solution Procedure

Our solution approach to the stochastic R&D portfolio optimization model involves a Lagrangian decomposition scheme. Note that model (20)-(24) is linked in scenarios through the nonanticipativity constraints (22)-(24). By subjecting these conditions to Lagrangian relaxation, we form the following Lagrangian

$$\begin{aligned} L(\mathbf{x}, \boldsymbol{\Upsilon}, \boldsymbol{k}, \boldsymbol{u}) = &\sum_{\omega \in \Omega} p^\omega \sum_t R_t u_t^\omega + \sum_{\omega \in \Omega} \sum_i \phi_i^\omega \left(\sum_{\omega' \in \Omega} p^{\omega'} \Upsilon_i^{\omega'} - \Upsilon_i^\omega \right) \\ &+ \sum_{\omega \in \Omega} \sum_{t \leq 5} \zeta_t^\omega \left(\sum_{\omega' \in \Omega} p^{\omega'} k_t^{\omega'} - k_t^\omega \right) + \sum_{\omega \in \Omega} \sum_{t \leq 5} \eta_t^\omega \left(\sum_{\omega' \in \Omega} p^{\omega'} u_t^{\omega'} - u_t^\omega \right) \end{aligned} \tag{41}$$

where ϕ_i^ω, ζ_t^ω, η_t^ω are the Lagrange multipliers. A major advantage of the described formulation of the nonanticipativity constraints is that when they are relaxed, the Lagrangian (41) can be decomposed by scenarios for given dual vectors ϕ, ζ, and η. Hence, the resulting Lagrangian can be expressed as

$$L(\mathbf{x}, \boldsymbol{\Upsilon}, \boldsymbol{k}, \boldsymbol{u}) = \sum_{\omega \in \Omega} L_\omega(\mathbf{x}^\omega, \boldsymbol{\Upsilon}^\omega, \boldsymbol{k}^\omega, \boldsymbol{u}^\omega) \tag{42}$$

The corresponding Lagrangian dual problem for problem (20)-(24) is then

$$\min_{\boldsymbol{\phi}, \boldsymbol{\zeta}, \boldsymbol{\eta}} \{ \mathcal{Z}(\boldsymbol{\phi}, \boldsymbol{\zeta}, \boldsymbol{\eta}) = \max \{ \sum_{\omega \in \Omega} L_\omega(\mathbf{x}^\omega, \boldsymbol{\Upsilon}^\omega, \boldsymbol{k}^\omega, \boldsymbol{u}^\omega) : (21) \} \tag{43}$$

Problem (43) is a nonsmooth convex minimization problem, and can be solved by subgradient optimization methods (Hiriart-Urruty and Lemarechal 1993). At each iteration of these methods, the solution of $\mathcal{Z}(\boldsymbol{\phi},\boldsymbol{\zeta},\boldsymbol{\eta})$ is required to obtain a subgradient. We note that $\mathcal{Z}(\boldsymbol{\phi},\boldsymbol{\zeta},\boldsymbol{\eta})$ is separable, and reduces to solving $|\Omega|$ problems of manageable size, each of which corresponds to a single scenario. Components of the subgradient vector Γ are then given by $\sum_{\omega'\in\Omega} p^{\omega'}\Upsilon_i^{\omega'} - \Upsilon_i^{\omega}$, $\sum_{\omega'\in\Omega} p^{\omega'}k_t^{\omega'} - k_t^{\omega}$, and $\sum_{\omega'\in\Omega} p^{\omega'}u_t^{\omega'} - u_t^{\omega}$ where Υ_i^{ω}, k_t^{ω} and u_t^{ω} are the corresponding optimal solutions to the scenario subproblems.

We let Γ^j represent the subgradient at iteration j, and propose a modified subgradient algorithm consisting of a combined step size rule. More specifically, we use a weighted combination of the subgradients from previous iterations in updating the dual variables, such that:

$$\hat{\Gamma}^j = \delta_0\Gamma^j + \delta_1\Gamma^{j-1} + \delta_2\Gamma^{j-2} \tag{44}$$

where the δ terms represent weights with $\delta_0+\delta_1+\delta_2=1$. Based on an experimental analysis of convergence rates, as it is the case for most subgradient algorithm implementations, we have determined that the best choices for these weights for the given problem are $\delta_0=0.7$, $\delta_1=\delta_2=0.15$.

Multiplier updates are then performed using the following step size rule:

$$\phi^{j+1} = \phi^j - \frac{\varphi(\bar{L}^j - \underline{L}^j)}{||\Gamma^j||}\hat{\Gamma}^j, \qquad \zeta^{j+1} = \zeta^j - \frac{\varphi(\bar{L}^j - \underline{L}^j)}{||\Gamma^j||}\}\hat{\Gamma}^j, \qquad \eta^{j+1} = \eta^j - \frac{\varphi(\bar{L}^j - \underline{L}^j)}{||\Gamma^j||}\}\hat{\Gamma}^j$$

where φ, $\varphi<2$, is a constant that can be modified during the algorithm, while $\bar{L}^j$ and $\underline{L}^j$ are upper and lower bounds on the Lagrangian at iteration j, respectively. The values to be used for φ were again determined through experimental analysis. Note that any Lagrangian dual solution is an upperbound for the original problem, which can be used in evaluating the value of a given feasible solution.

Despite the improvements in convergence rates through the parameter settings above, the subgradient algorithm implementation is still not efficient enough for quick evaluations of the large number of policy environments and input configurations that we have considered as part of our analysis in this paper. However, further improvement of the solution procedure is possible by establishing the following result about the structure of the optimal investment decisions for the given piecewise linear returns functions.

PROPOSITION 2. *If* $\lambda_i^{n\omega,*}$ *represent the optimal values for variables* $\lambda_i^{n\omega}$*, then* $\lambda_i^{n\omega,*}\in\{0,1\}$ *for all* n*,* i *and* ω*, i.e. the optimal investment decision for each technology category* i *corresponds to a vertex value in the corresponding piecewise linear returns function.*

Proof: The result follows from a marginal analysis. Consider equation (14) as defined for each scenario $\omega\in\Omega$. Given that maximization of the utility in each period implies the maximization of the net output y_t^ω for that period, it is optimal to increase investment by Δ_i units as long as $E_\omega[\Delta_i^{y,\omega}] \geq \kappa\frac{\Delta_i}{5}$, where $\Delta_i^{y,\omega}$ is the change in the net output value of scenario ω for a Δ_i unit increase in investment for technology category i.

By definition, the marginal returns and costs are the same in the range $\Upsilon_i \in [v_i^n, v_i^{n+1}]$ for all n due to the linear relationships between investment levels and α_i^ω. Suppose for some i, n and ω, $0 < \lambda_i^{n\omega,*} < 1$, i.e. the optimal investment is not a vertex value implying that $v_i^n < \bar{v}_i < v_i^{n+1}$, where $\Upsilon_i^* = \bar{v}_i$. Assuming without loss of generality that the expected returns are increasing between vertices n and $n+1$, the optimality conditions imply that $E_\omega[\Delta_i^{y,\omega}] \geq \kappa\frac{\Delta_i}{5}$ in the range $\Upsilon_i \in [v_i^n, \bar{v}_i]$. On the other hand, this should also hold for the range $\Upsilon_i \in [\bar{v}_i, v_i^{n+1}]$ due to the constancy of marginal returns between vertices n and $n+1$. Hence, it is possible to increase social utility by increasing the investment level to the value at vertex $n+1$, which is a contradiction implying that $\bar{v}_i$ can not be optimal. This would require $\lambda_i^{n\omega,*} \in \{0,1\}$ for all n, i and ω.$\square$

Hence, it is possible to implement an implicit enumeration procedure for the investment levels at the vertices of the piecewise linear returns functions and only solve for the optimal abatement policy at those implicitly enumerated investments levels. Implementation of this procedure improves the overall solution time significantly as the subgradient iterations are only implemented over the variables u and k for given investment levels.