Optimal Climate Change Policy: R&D Investments and Abatement under Uncertainty

Erin Baker
Department of Mechanical and Industrial Engineering, College of Engineering, University of Massachusetts, Amherst, MA 01003, edbaker@ecs.umass.edu

Senay Solak
Department of Finance and Operations Management, Isenberg School of Management, University of Massachusetts, Amherst, MA 01003, solak@som.umass.edu

Human induced climate change is a major public policy problem today. One vexing problem faced by policy makers is how to allocate research budgets across a variety of energy technologies, in order to reduce the future costs of controlling climate change. In this paper we apply a multi-model approach, implementing probabilistic data derived from expert elicitations into a stochastic programming version of a dynamic Integrated Assessment Model, in order to arrive at insights about the optimal government-funded R&D portfolio. We focus on electricity technologies with a significant chance of a breakthrough. We find that the optimal investment is fairly robust to different specifications of climate uncertainty, to different policy environments, and to assumptions about the opportunity cost of investing. We also identify results about the cost of investing nonoptimally under different risk cases, and reach important conclusions about the role R&D plays in different types of policy environments.

Key words: R&D portfolio, energy technology, climate change, stochastic programming, public policy, DICE

1. Introduction

Climate change is one of the biggest public policy problems facing the world right now. 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. As in any stochastic decision framework, explicit inclusion of uncertainty and learning into the climate change problem can have significant impacts on optimal policy. It has been shown theoretically, however, that these impacts are mostly ambiguous, i.e. can not be categorized as increasing or decreasing with respect to varying degrees of risk (Baker and Shittu 2008, Baker 2009). Specifically, this implies that certain increases in risk will lead to optimally more abatement or R&D investment, while other increases in risk will result in optimally less abatement or R&D investment. Therefore, in order to understand how the optimal near term policy should be defined under uncertainty, we need to incorporate specific valid probabilistic characterizations of relevant stochastic parameters into the problem framework, and analyze the corresponding complex decision problem using this currently available stochastic information. Our work, culminating in this paper, is designed to achieve this by going beyond the theoretical and providing a framework to collect and implement probabilistic data into climate change policy analysis through tractable stochastic optimization, with a particular focus on technology R&D policy.

There are two key, near-term avenues for response to climate change. The first, most direct avenue is *abatement,* that is to reduce emissions of the greenhouse gases (GHG) 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 reduction will be less costly in the future.

These two avenues, however, are not independent. A given emissions path influences the set of technologies we would like to have in the economy; and the set of technologies actually available influences the optimal level of emissions reductions. We explicitly recognize and model this interdependency as part of our analysis in this paper. Specifically, we simultaneously determine the optimal investment in a portfolio of technology R&D projects and the optimal emissions path so that the expected societal costs of climate change are minimized.

Determination of an optimal energy technology R&D policy is hard for a number of reasons. But one of the most salient reasons is uncertainty. There is deep uncertainty around climate change damages. In particular, it is not known what the marginal damages caused by one more ton of GHG emissions are. There is also uncertainty about technical change. We don't know which R&D projects will be successful at what level.

As noted above, there has been 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). On the other end of the spectrum from the theoretical analysis is a body of work based on technologically detailed Integrated Assessment Models (IAMs). 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 of different environmental goals and different assumptions about technologies. For example, a number of models have considered the impacts on costs of having or not having carbon capture and storage (CCS) available (Clarke et al. 2008, 2009). While these top-down analyses through IAMs 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. Yet, a recent study by the National Academy suggests that uncertainty be explicitly included in U.S. Department of Energy (DOE) decisions about investments into R&D (NRC 2007). In this work we build on the technological detail of an IAM, and go beyond sensitivity analysis to specifically incorporate uncertainty into such decision making. To the best of our knowledge, this is the first such integrated approach that combines elicited probabilistic data and stochastic optimization with detailed assessment models.

Specifically, we implement data on technological change based on expert elicitations into a well-known IAM, Dynamic Integrated Model of Climate and the Economy (DICE) (Nordhaus 1993). We extend the original DICE model to include uncertainty and learning about technology and climate damages using stochastic programming. Thus, we provide a study that goes beyond theoretical analysis on the one hand, and data-based sensitivity analysis on the other hand. This requires a multi-step, multi-model process. In the next section we describe the multiple steps and models that form the background of the current paper. In Section 3 we describe our current model and the methodology we use to solve it.

In Sections 4 and 5 we use our resulting model to explore the role of R&D in a number of different policy environments and risk cases. We follow the lead of Nordhaus (2008) in modeling a set of alternative policy environments, including a baseline no-controls scenario, a scenario limiting temperature change to no more than 2°C, a scenario based on the spirit of the Kyoto Protocol, and two ambitious scenarios based on proposals by Gore (Gore 2007) and the Stern Review (Stern 2007),

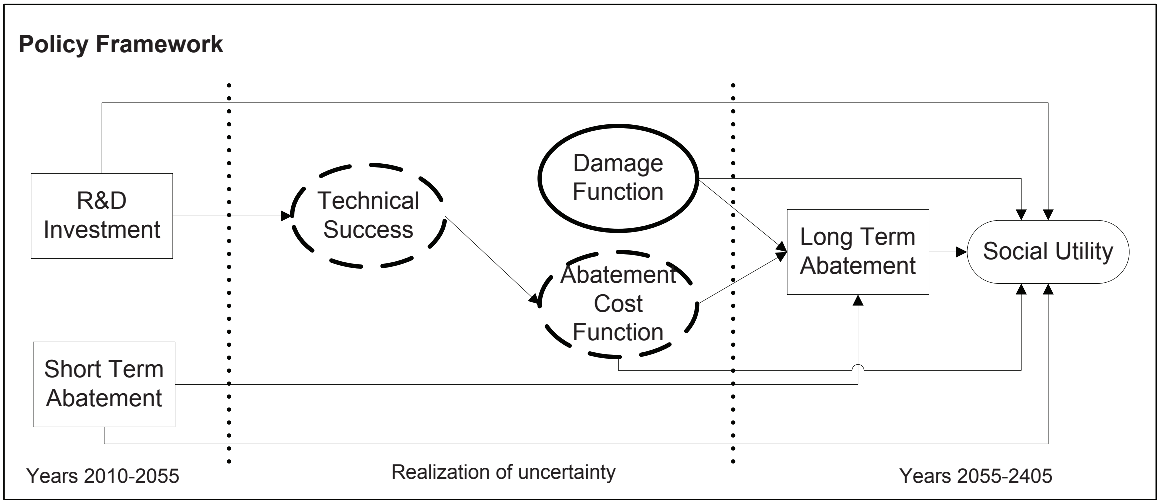

Figure 1 **Influence diagram for the decision process to maximize social utility in the presence of climate change and technological uncertainty**

along with the optimal policy chosen by DICE. We also consider a set of risk cases, representing different probability distributions over climate damages. We compare the optimal R&D investment in the different environments and cases, and compare the value of the policy environments with and without R&D. We find, with one exception, that the optimal R&D investment is quite robust across policy environments and risk cases. The one exception is that R&D investment is much higher if it is chosen under a low discount rate, consistent with the Stern Report.

Overall, we also find an asymmetry in the value of R&D investments around the optimal investment level. More specifically, the marginal increase in expected costs for investment amounts that are above the optimal investment level is lower than the marginal cost increase for amounts less than the optimal. This implies that a conservative investment policy is not preferable.

We conclude with a discussion on policy insights in Section 6.

2. Background: A Multi-step Multi-model Process

In this section, we first describe the decision problem that we are interested in, and then discuss the multi-step multi-model process used to generate the necessary inputs for the comprehensive analysis in this paper.

Figure 1 is an influence diagram representing a very simplified version of our problem. The near term decisions involve how much to invest in which technologies and the level of near term abatement. The arrow from the R&D investment decision node to the uncertainty node over technological success means that the probability distribution over technological success depends on the projects that are funded. Hence, technical success is defined as an endogenous stochastic parameter in the decision model, indicated by a dotted oval. The next relationship is a deterministic mapping between technical success and the abatement cost functions, where abatement is defined as a reduction in emissions below a business-as-usual level. It is important to note that technology impacts the entire cost curve, not just a point estimate of abatement costs. The second stochastic characterization is an exogenous one, indicated by a regular oval, and involves the damages from climate change, which are also represented as a function, dependent on the stock of emissions in the atmosphere. It is assumed that our second stage decision – long term abatement policy – will be determined after information about abatement costs and the damages becomes available. The overall objective is to maximize expected total social utility. Notice that, in the absence of a corner point, abatement will be optimized where the marginal cost of abatement is just equal to the marginal damages avoided by abatement. Thus, an important unit of analysis is the Marginal Abatement Cost Curve (MAC), which is the curve that reflects the cost of reducing emissions by an additional ton. Surrounding all these decisions is the overall policy framework.

The first challenge in the analysis is to determine how R&D investments affect the probability of success over technologies. Since we are interested in identifying an optimal portfolio, we need to differentiate between different energy technologies and projects. Moreover, we are particularly interested in breakthroughs – game-changing technical advances that may be quite different from anything that has gone before. Therefore, no past data is available to be used in characterizing stochastic returns for candidate technologies. The only way to gather such information, in fact, is to derive subjective probabilities based on expert judgement.

Thus, the *first step* of the multi-model process, which we outline in Figure 2, was to perform expert elicitations with scientists and engineers familiar with the specific technologies. As part of the elicitations, we developed definitions of success for each technology at multiple levels. As the

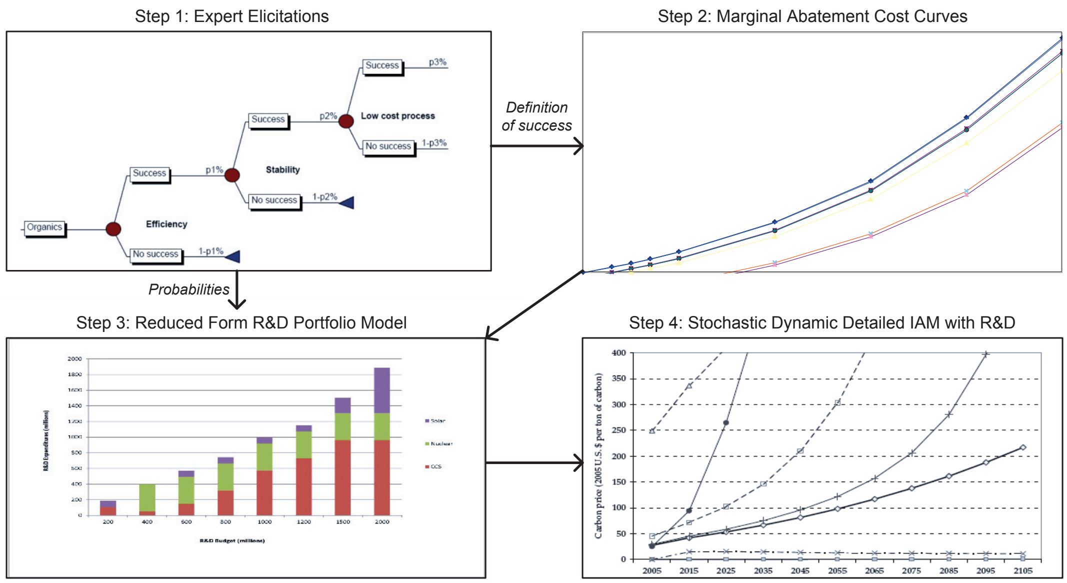

Figure 2 **Overview of the multi-step multi-model process to determine optimal investment and abatement policies in response to climate change**

second step, we used these definitions of success in a technologically detailed integrated assessment model, the Global Change Assessment Model (GCAM), to derive MACs. We then combined the MACs with the probabilities we elicited to get probability distributions over MACs, conditional on R&D funding levels. This data was then used in the *third step* as part of a reduced-form R&D model to gain some insights about the characteristics of an optimal portfolio. In this study, we use these insights to implement the collected data into a dynamic, detailed IAM, which we significantly extend to include uncertainty and learning within a stochastic optimization framework, and obtain policy-based results. This represents the *fourth step* in the process. In the rest of this section we start by discussing our choice of technologies, and then briefly describe each of the previous steps in the process.

2.1. The Technologies

Our analysis considers investment options in three key technology areas: solar photovoltaics, carbon capture and storage, and nuclear fission. Here we discuss why we chose these particular technologies to focus on. In the next section we briefly describe how we quantify their impact on abatement.

There are many technologies and technology categories with some relevance to climate change.

The DOE invests in energy R&D primarily in ten large categories: hydrogen, biomass, geothermal, solar, hydro, nuclear, wind, energy conservation, clean coal, and CCS. Each of these large categories can be divided into a number of sub-categories. Here we note that our work has two objectives: to directly inform DOE investment policy, and to provide and test a framework that can be expanded in the future to include more technologies as necessary. Thus, we focused our attention on a subset of technologies, based on three key criteria that we discuss here.

First, we concentrate on electricity generation technologies. Electricity is the single largest creator of carbon emissions, and the sector is predicted to increase in importance under climate change. This largely removes hydrogen, biomass, and energy conservation from interest.

Second, we are interested in technologies in which the possibility of technological breakthroughs plays a significant role. That is, our primary interest is on the *research* in R&D. This is because incremental technical change is more amenable to deterministic modeling techniques such as "learning by doing". While these incremental changes are important, our main goal is to attack the most difficult part of the problem, the part that is fraught with uncertainty. Some of the relevant energy technologies are not as subject to breakthroughs, but appear amenable to more straightforward engineering. The key technology in this category is wind. The frontier for wind energy is large, offshore turbines. While these will require a number of engineering advances, they probably don't require any scientific breakthroughs. The key advances are related to increasing scale and reliability. Three other technologies with similar characteristics are hydro, solar thermal, and clean coal. These kinds of incremental engineering advances are typically included in the baseline of most energy-economics models, which include the DICE IAM we utilize as part of our analysis.

Third, we are interested in technologies that have the potential to play a very significant role in the response to climate change. One aspect that determines how large a role a technology will play is its resource base. Lewis and Nocera (2006) have pointed out in their analysis that solar, nuclear power, and carbon capture and storage are the three technologies with sufficient resources to provide the carbon-neutral energy needed to address the climate change problem. On the other end of the spectrum from these technologies is hydro with a resource base estimated to be less than

1 TW, with nearly 0.6TW already in place. We also place geothermal in this category, although the total resource base is quite controversial. In fact, this suggests that the primary direction for initial research on this technology is into the economically and politically exploitable resource base.

Moreover, we base our choice of technologies on the results from large, technologically detailed IAMs, which integrate economic and climate models. The Climate Change Science Program (CCSP) has released a report in which the results of three important global models are compared (Clarke et al. 2007). The results of the three models indicate that electricity will grow in importance in the energy sector and will be about 25% of all primary energy production by 2100. It is also noted that the combination of renewables, CCS, and nuclear will make up between 93% and 100% of electricity generation. In fact, these technologies go beyond the electricity sector, making up a total of 39 - 68% of total primary energy production. The renewable technologies vary by model and scenario, but generally include solar photovoltaics (PV), solar thermal, wind, hydro, and geothermal. Given the discussion above for solar thermal, wind, hydro and geothermal, solar PV is left as the most interesting renewable.

Hence, due to these issues, we focus on R&D investment options in three key technologies: nuclear fission, CCS, and solar PV. Each of the three key technology categories contains multiple research areas in which R&D investments can be made. These research areas are listed in the project column in Tables 1-3. In Table 1, *inorganic solar cells* refer to a search for better inorganic semiconductors, to replace silicon and other less promising but well-studied alternatives, while *third generation technologies* include highly efficient technologies involving new cell architectures, quantum dots and multi-junction cells. A brief overview of all the technologies can be found in Baker and Solak (2010) and detailed descriptions in Baker et al. (2008), Baker et al. (2009a) and Baker et al. (2009b).

2.2. Deriving Random Marginal Abatement Cost Curves

We performed three sets of elicitations on the three technology categories. Elicitation results are summarized in Tables 1-3. The tables show the amount of R&D funding required for each project

Technology	Project	NPV of Funding (000,000)	Probabiltiy of success	Alpha	Shift
Solar	*Organic Solar Cells*	$116	0.0%	0.050	0.017
			13.0%	0.022	0.007
		$830	3.9%	0.050	0.017
			24.8%	0.022	0.007
	Inorganic Solar Cells	$39	26.7%	0.022	0.007
		$77	44.3%		
	3rd Generation Technologies	$386	2.0%	0.050	0.017

Table 1 **Summary of expert elicitation results for the solar technology.**

Technology	Project	NPV of Funding (000,000)	Probabiltiy of success	Alpha	Shift
CCS	*Pre-combustion Carbon Capture*	$39	2.7%	0.347	0.004
		$154	11.0%		
		$386	22.3%		
	Chemical Looping	$19	8.0%	0.380	0.020
		$38	29.5%		
		$56	42.0%		
	Post-combustion Carbon Capture	$52	59.0%	0.319	-0.008
		$22	70.0%		
		$519	78.5%		

Table 2 **Summary of expert elicitation results for the CCS technology.**

Technology	Project	NPV of Funding (000,000)	Probabiltiy of success	Alpha	Shift
Nuclear	*Light Water Reactors (LWR)*	$173	21.3%	0.325	0.118
		$260	33.8%		
		$346	60.0%		
	High Temperature Reactors (HTR)	$772	0.3%	0.327	0.111
			0.9%	0.111	0.028
		$1,544	17.0%	0.327	0.111
			9.2%	0.111	0.028
		$3,089	30.2%	0.327	0.111
			10.1%	0.111	0.028
	Fast Burner Reactors (FR)	$1,158	0.1%	0.332	0.115
			7.4%	0.115	0.029
		$4,633	0.5%	0.332	0.115
			32.0%	0.115	0.029
		$15,443	16.3%	0.332	0.115
			43.8%	0.115	0.029

Table 3 **Summary of expert elicitation results for the nuclear technology.**

and the probability of success. The last two columns represent the parameters we used to represent the impact of technology on the abatement cost curve, which we discuss later in this subsection.

In Figure 3 we show a selection of the MACs that we derived using GCAM (Brenkert et al. 2003, Edmonds et al. 2005). The baseline MAC is based on assumptions defined by Clarke et al.

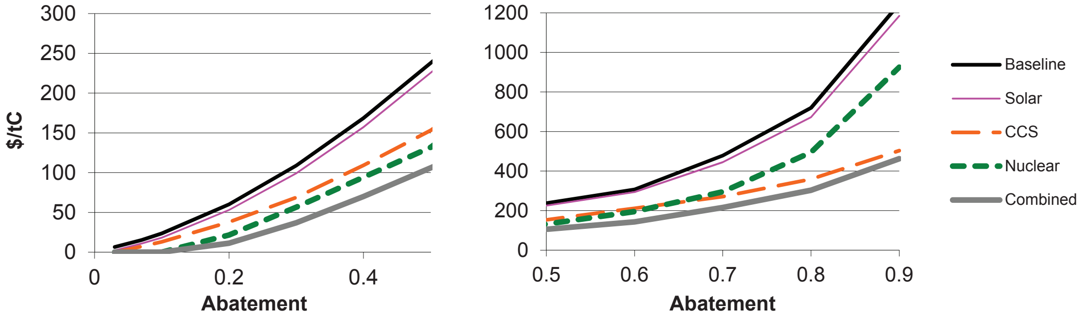

Figure 3 **Representative MACs**

(2008). We show the impact of the LWR project, labeled Nuclear, chemical looping, labeled CCS, and inorganic solar, labeled Solar. Finally, we show the MAC that is derived if all three of the technologies were successful.

Here we note that solar PV is an intermittent technology – it cannot be turned on or off at will, but rather works precisely when the sun shines. This causes potential problems for the electricity grid, including potentially large investments in rarely-used generation capacity and/or potential instability. In our model, our baseline assumption is based on the baseline assumption in GCAM in which PV with battery storage competes with PV without storage. As the penetration of PV without battery storage into the electricity grid approaches a share limit of 20%, each additional kW of PV requires a larger amount of gas-fired backup generator, until each kW of PV requires one kW of gas-fired backup (and the ratio stays constant thereafter). PV with battery storage generates electricity without any backup requirement, but has an additional cost to account for the batteries.

Note that CCS and Nuclear have very different impacts on the MAC. CCS primarily pivots the cost curve down. This is because CCS will have no effect in the absence of a carbon policy, but will have a significant effect at very high levels of abatement. CCS allows for the continued use of coal-fired generation, which currently makes up more than 50% percent of all generation in the U.S. On the other hand, nuclear shifts the curve to the right, with a much smaller pivot effect. This is because an improvement in nuclear will reduce carbon emissions even in the absence of a carbon

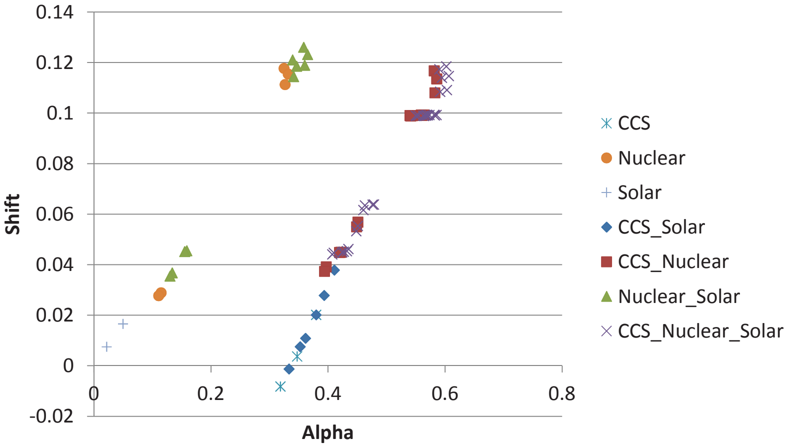

Figure 4 **The MAC shift and pivot parameter values of all technology combinations.**

policy. On the other hand, GCAM uses past data on how nuclear has competed in the economy to predict future results. Thus, nuclear, which has been slow to diffuse in relation to its costs, does not diffuse enough to have a pivoting effect at high levels of abatement. Solar has a similar qualitative impact as nuclear, although its overall impact is small, even when grid integration problems are relaxed. Given this, we consider nuclear and solar as a single category in our model in this paper.

Based on the results from GCAM, we estimated parameters to represent the pivoting and shifting effects of the different technologies on the MAC. These parameters are shown in Figure 4, with alpha measuring the pivot. While the pivots of a combination of technologies are determined to be multiplicative, the shifts were estimated separately for each combination. Baker and Solak (2010) provides a detailed discussion of these parameterizations.

2.3. Reduced Form R&D Model

The next step in our analysis was the implementation of our data into a reduced form R&D portfolio model in order to derive parametric input characterizations for our advanced integrated policy model in this paper. This reduced form model is described in detail in Baker and Solak (2010), where we focused our attention on the individual R&D projects, and reduced the economy into two periods and a single equation, calibrated to the DICE IAM.

This model provided two key results that underline our analysis in this paper. First, the composition of the optimal portfolio was robust to risk in climate damages, where risk is defined in terms of a mean-preserving spread (MPS) (Rothschild and Stiglitz 1970). That is, at any given budget level, the composition of the optimal portfolio did not change with changes in the damage risk. In fact, it did not even change with changes in the *level* of damages. While we show theoretically that this is not the case for any general instance of the problem, this robustness property holds for the elicited data that defines the current state of knowledge on technological success. This is because projects are quite differentiated based on the elicited data, with some projects having low investments, high success probabilities, and large impacts when successful, while other projects are generally weaker. This result is very useful to our current analysis. It means that we do not have to explicitly model the individual technology projects. More specifically, for each budget level we can define a probability distribution over the pivot and shift parameters, and then use these to derive returns functions for different technologies. This step is further described in Section 3.2.2.

Second, we found that while the *composition* was robust for a given budget, the optimal amount of R&D funding does change with climate risk in a non-monotonic way. Specifically, when we model an increase in risk as a mean-preserving-spread that stretches out the tail, we find that the optimal amount of R&D funding first increases, and then decreases in risk. The reason has to do with the interplay between technology investment and long term abatement. Mean damages in DICE are equivalent to a 1.1% loss in GDP given a 2°C increase in global mean temperature, and medium term abatement (2050 - 2100) ranges between about $25-45\%$ in the absence of technical change. When technology is successful, however, abatement increases radically, to between $45-75\%$. This leads to lower damages, but actually leads to higher total expenditures on abatement in the deterministic case. In a medium-risk case, with a $1/3$ chance that damages are 3 times higher, we find a slightly different pattern. Abatement increases with technical change, thus lowering damages. But, the total cost of abatement is also lower after technical change, thus maximizing the benefit of technical change. In this case, we find that the optimal investment in technology is high. Finally, when damages are very high, it is optimal to abate fully regardless of the outcome of technical change. Thus, the presence of technology does not have an environmental-side effect as damages are not changed by technology. Only the cost of abatement is reduced. This leads to a lower value for technical change in this scenario, and a lower optimal investment. This result supports the motivation for our current model, in which we use a stochastic dynamic optimization model in order to capture the complex interactions between climate change damages, abatement, and the impact of R&D.

3. Integrated R&D and Abatement Policy Optimization Model

In this section we present our data-based stochastic optimization model. We start by describing the basic DICE model and then discuss how we extend this basic model to include uncertainty and learning, specifically how we incorporate our probabilistic data on technical change into the model. We then describe our stochastic optimization framework and solution concepts.

3.1. 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, output is divided between consumption and investment in new capital, consistent with the standard optimal growth framework. DICE adds to this framework by allowing for emissions of GHG into the atmosphere as part of the production process. The accumulation of the GHG affects welfare by increasing global temperatures, which, in turn, reduce production. In order to mitigate this effect, an abatement level can be chosen each period, which reduces emissions below what would otherwise occur for a given production level. While abatement has obvious benefits, it is costly, reducing the amount of output available for consumption or investment in every period. The objective of the model is to maximize the discounted sum of social utility over time, where utility is based on consumption. The optimal abatement path reflects a balance between benefits and costs. The general formulation of the DICE model is given in Nordhaus (2008). Here we highlight a key equation from the model which shows how the cost of abatement and the damages from climate change impact output y_t in each period t:

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

where y_t^g is unadjusted output, $c(\mu_t)$ is the cost of abatement μ_t, i.e. the fraction of emissions reduced below what would otherwise take place, and $D(\tau_t) = 1 + \pi\tau_t^2$ is the damages from climate change resulting from temperature τ_t, where $\pi = 0.00284$ in the original implementation of the DICE model. We will take π to be a stochastic parameter in our model, modeling the uncertainty in climate change damages.

3.2. Implementing the Data

3.2.1. Modeling Technical Change As discussed in Section 2, we estimated the impact of each individual technology on the MAC based on a parameterization of the shifting and pivoting effects. Since nuclear and solar have similar impacts, we have combined them into one category. In this section we describe how we integrate this data into our model.

First, we consider the investment decisions for energy technologies, assuming that R&D investment will take place over the next 50-year period at a constant rate. In the base DICE model, output is divided between consumption c_t, and investment in traditional capital l_t in each period. We extend this relationship to include R&D investment for periods $t \leq 5$ as follows:

$$y_t = c_t + l_t + \kappa \sum_{i=1}^{2} \Upsilon_i / 5 \qquad \forall t \leq 5 \tag{2}$$

where Υ_i is the investment in each technology area, κ is the opportunity cost parameter, and the indices 1 and 2 correspond to CCS and solar/nuclear technology categories, respectively.

Second, we consider how technical change impacts the cost of abatement. The cost of abatement in DICE is represented by

$$c(\mu_t) = [1 - P_t B_t \mu_t^{\theta}] \qquad \forall t \tag{3}$$

where $\theta = 2.8$ and $P_t B_t$ 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 “markup” in the costs related to participation. 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.

We measure the returns from technology investments by the pivot parameters α_i, $i = 1, 2$ and the shift parameter h, which affect the abatement cost function and thus the net economic production equation (1) in DICE. The probabilistic estimates of these parameters for different investments levels are given in Section 3.2.2. Given these, the cost of abatement after technical change can be expressed as follows:

$$c\left(\mu_t, \alpha, h\right) = \left(1 - \alpha_1\right)\left(1 - \alpha_2\right)\left[c\left(\mu_t\right) - hc\left(0.5\right)\mu_t\right] \qquad \forall t \tag{4}$$

where the shift, h, is anchored on 50% abatement to make it portable. Note that a one-to-one mapping can be defined between each pair of possible α_i values and the shift parameter h, allowing the representation of this modified cost function as a function of α_i values only. We use this relationship as part of our integrated R&D and policy optimization model, but represent it through piecewise linear approximations as discussed in Section 3.3.

3.2.2. Modeling 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 abatement costs which is a function of the R&D investment decisions.

We model the stochasticity in climate change damages through a probabilistic characterization of the parameter π in the damage equation $D\left(\tau_t\right) = 1 + \pi {\tau_t}^2$. 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 4.2.

For the uncertainty in technical change, we develop a probabilistic mapping from investment decisions to the technical change parameters α_i, $i = 1, 2$. Note that the reduced-form R&D model that we describe in Section 2.3 gives us optimal portfolios for a set of investment levels. These levels are defined by considering a set of possible R&D budgets, and then using the reduced-form model to solve for the optimal portfolio at each budget level. Due to the discrete nature of project selection decisions in the portfolio problem, and the resulting 0-1 knapsack structure, the actual investment amounts are typically less than the considered budget levels. The actual set of investment levels determined for CCS and solar/nuclear are shown in the budget rows in Tables 4

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 4 Piecewise linear returns functions for CCS.

and 5, respectively. Using this information, it is possible to derive probabilities that are dependent on investment decisions for different values of α_i for each technology category. However, such endogenous probabilities typically are not amenable to stochastic optimization methods, specifically stochastic programming. In order to make the data amenable for stochastic programming, 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 assign an α_i value to each R&D investment level.

The "returns functions" are derived by first considering all possible success outcomes for technologies as defined in Tables 1-3, and then by identifying the corresponding α_i values for each investment level based on the composition of the optimal portfolio at that investment level as determined by our previous model. Hence, given the elicited data, it is possible to completely characterize probability distributions through enumerating all possible returns functions and calculating the corresponding probabilities. However, these enumerations result in a large number of possible functions, especially for the solar/nuclear category, implying an intractable optimization model. 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 4 and 5 show the approximated returns functions for the CCS and solar/nuclear technology categories, respectively. For example, in Table 5 the third row of the data represents a "returns

Budget($mil)	77	346	423	539	4014	8975	20171	Probability
α_2	0	0	0	0	0	0	0	0.03
	0.022	0	0.022	0.022	0.022	0.022	0.022	0.13
	0	0	0	0	0	0.115	0.115	0.05
	0	0	0	0	0.134	0.134	0.134	0.02
	0.022	0.325	0.34	0.34	0.34	0.34	0.34	0.17
	0	0.325	0.325	0.34	0.34	0.34	0.34	0.04
	0	0.325	0.325	0.325	0.325	0.325	0.325	0.17
	0	0	0	0	0.327	0.327	0.327	0.10
	0.022	0	0.022	0.022	0.342	0.342	0.342	0.11
	0	0	0	0.022	0.342	0.342	0.342	0.03
	0	0	0	0	0	0	0.347	0.06
	0	0.325	0.325	0.325	0.325	0.325	0.359	0.06
	0.022	0.325	0.34	0.34	0.342	0.342	0.342	0.03

Table 5 **Piecewise linear returns functions for solar/nuclear.**

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 6 **Comparison of the estimated and actual data for CCS.**

function" that has α_2 equal to zero for an investment less than \$8,975 million, and 0.115 for a larger investment. The probability that this particular function is realized is 0.05. 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 6 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 budget 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.

3.3. Stochastic Optimization Model

In this section we discuss how we implemented a two-stage stochastic programming model using DICE as the basis. Given the representation of the uncertainty in technical change and climate change damages as described in Section 3.2.2, the stochastic optimization problem involves the determination of an optimal portfolio of technology investments and an abatement policy such that the expected total social utility is maximized.

In dealing with this stochastic optimization problem, stochastic programming is a natural approach as the interactions in the DICE model form a complex structure that prevents the problem from being amenable to other methods such as dynamic programming.

The two-stage stochastic programming model is constructed by considering the set of scenarios formed by combinations of the possible realizations of the stochastic parameters, which consist of parameters defining the "returns functions" and the damage uncertainty. Individual scenario problems are then linked through the nonanticipativity constraints involving the investment, abatement and period utility decisions for the first 50 year period, i.e. Υ_i, μ_t, and u_t, respectively. It must be noted here that the set of decision variables at each period in the DICE model involves a large number of other variables. However, the following result shows the sufficiency of nonanticipativity in the three decisions above for the overall model:

Proposition 1. *For the extended DICE model with investment decisions, let* Υ_i^ω, μ_t^ω, u_t^ω, *and* $\mathbf{x}_t^\omega$ *represent the decision variable values for scenarios* $\omega \in \Omega$, *where* $\mathbf{x}_t^\omega$ *is the vector of all other variables.*

For any $\omega, \omega' \in \Omega$, *if* $\Upsilon_i^\omega = \Upsilon_i^{\omega'}$, $\mu_t^\omega = \mu_t^{\omega'}$, *and* $u_t^\omega = u_t^{\omega'}$, *then* $\mathbf{x}_t^\omega = \mathbf{x}_t^{\omega'}$.

Proof: The result can be established by considering all feasible relationships between the decision variables, and recognizing that the objective involves a maximization.□

Letting p^ω represent the probability of scenario ω, the stochastic DICE-based portfolio model can be expressed in general form as follows:

$$\max \sum_{\omega \in \Omega} p^\omega F^\omega(\mathbf{x}, \Upsilon, \mu, u) \tag{5}$$

$$\text{s.t. } G_m^\omega(\mathbf{x}, \Upsilon, \mu, u) \le b_m \qquad \forall m, \omega \tag{6}$$

$$\Upsilon_i^\omega - \sum_{\omega' \in \Omega} p^{\omega'} \Upsilon_i^{\omega'} = 0 \qquad \forall i, \omega \tag{7}$$

$$\mu_t^\omega - \sum_{\omega' \in \Omega} p^{\omega'} \mu_t^{\omega'} = 0 \qquad \forall t \le 5, \omega \tag{8}$$

$$u_t^\omega - \sum_{\omega' \in \Omega} p^{\omega'} u_t^{\omega'} = 0 \qquad \forall t \le 5, \omega \tag{9}$$

In this formulation, the function $F^\omega(\mathbf{x}, \Upsilon, \mu, u)$ represents the social utility function for a given scenario ω, while constraints (6) define the corresponding set of constraints for the scenario. These constraints, each of which is indicated by $m = 1, \ldots, M$, involve the standard economic relationships described in DICE, as well as the extended relationships modeled by equations (2) and (4). Constraints (7)-(9) are the nonanticipativity constraints ensuring that decisions in the first 50 years are the same for all scenarios, as established by Lemma 1. Notice that the structure used in the formulation of the nonanticipativity constraints account for the scenario probabilities, and prevent the ill-conditioning in the Lagrangian dual, which we use in Section 3.4. The possibility of ill-conditioning in other types of formulations is discussed by Louveaux and Schultz (2003).

A major complication in the described model is the existence of several nonlinear relationships in the set of constraints involving multiple variables, and the implied potential nonconvexities. While experimental analysis has shown that the original DICE model is likely convex, a theoretical analysis based on this conjecture is also available in Solak and Baker (2010). Hence, for the convexity of our stochastic model, it suffices to ensure that the extended relationships we model, namely equations (2) and (4), are also convex.

While the above is not an issue for the linear relationship in (2), equation (4) is highly nonconvex due to the multilinear terms involving multiplication of the pivot and shift values which are directly related to the decision variables of how much to invest in R&D. To deal with this, we have estimated tight linear approximations to the actual functions. More specifically, we used our data on the α_i and h values, and estimated the two following quantities:

$$(1 - \alpha_1)(1 - \alpha_2) \approx 1 - 0.8\alpha_1 - 0.92\alpha_2 \tag{10}$$

$$(1-\alpha_1)(1-\alpha_2)h \approx 0.02-0.06\alpha_1+0.14\alpha_2 \tag{11}$$

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

$$y_t = \frac{\left[1-\left((1-0.8\alpha_1-0.92\alpha_2)c(\mu_t)-(0.02-0.06\alpha_1+0.14\alpha_2)c(0.5)\mu_t\right)\right]}{D(\tau_t)} y_t^g \tag{12}$$

It is shown by Solak and Baker (2010) 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. This can be done by establishing that the equality relationship in (12) can also be stated as an inequality and then by analyzing the properties of the right hand side in (12).

The second complexity in our model is the representation of the stochastic returns functions. By the nature of their derivation which is based on a discrete set of investment levels, each of these functions is piecewise linear. The uncertainty in these functions is modeled using a set of parameters corresponding to the vertices used in the piecewise linear curves. More specifically, we define new variables $\lambda_i^k \geq 0$ for $i = 1, 2$ and $k = 0, \ldots, K_i$, where K_i is the number of vertices used to represent the return functions for technology category i, and then include the following constraints in the modified stochastic DICE model:

$$\Upsilon_i^\omega = \sum_k^{K_i} v_i^k \lambda_i^{\omega k} \qquad \forall i, \omega \tag{13}$$

$$\alpha_i^\omega = \sum_k^{K_i} \hat{\alpha}_i^\omega(k) \lambda_i^{\omega k} \qquad \forall i, \omega \tag{14}$$

$$\sum_{k=1}^{K_i} \lambda_i^{\omega k} = 1 \qquad \forall i, \omega \tag{15}$$

where v_i^k is the investment value for the kth vertex. The stochastic parameter $\hat{\alpha}_i^\omega(k)$ in these constraints corresponds to the value of the return parameter α_i^ω at the kth vertex of the return function. Note that we must require that at most two adjacent $\lambda_i^{\omega k}$ can be nonzero for each i and ω. This condition ensures that corresponding values of Υ_i^ω and α_i^ω lie on one of the straight line segments of the return function. However, this condition will be satisfied regardless due to the convexity of the original model.

3.4. Solution Methodology

Model (5)-(9) is a two-stage stochastic nonlinear problem that can solved be through decomposition methods, as the size of the scenario set does not allow for the direct solution of the deterministic equivalent of the problem. This is still the case despite the reductions in the scenario set through the approximations of the returns functions. Hence, we use a Lagrangian decomposition based procedure to solve the corresponding problem, which is similar to the method described by Caroe and Schultz (1999).

Note that model (5)-(9) is linked in scenarios through the nonanticipativity constraints (7)-(9). By subjecting these conditions to Lagrangian relaxation, we form the following Lagrangian

$$\begin{aligned} L(\mathbf{x},\Upsilon,\mu,u) = & \sum_{\omega\in\Omega} p^{\omega} F^{\omega}(\mathbf{x},\Upsilon,\mu,u) && (16)\\ & + \sum_{\omega\in\Omega}\sum_{i} \phi_i^{\omega}\left(\sum_{\omega'\in\Omega} p^{\omega'}\Upsilon_i^{\omega'} - \Upsilon_i^{\omega}\right) && (17)\\ & + \sum_{\omega\in\Omega}\sum_{t\leq 5} \tau_t^{\omega}\left(\sum_{\omega'\in\Omega} p^{\omega'}\mu_t^{\omega'} - \mu_t^{\omega}\right) && (18)\\ & + \sum_{\omega\in\Omega}\sum_{t\leq 5} \gamma_t^{\omega}\left(\sum_{\omega'\in\Omega} p^{\omega'}u_t^{\omega'} - u_t^{\omega}\right) && (19) \end{aligned}$$

where ϕ_i^{ω}, τ_t^{ω}, γ_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 (19) can be decomposed by scenarios for given dual vectors ϕ, τ, and γ. Hence, the resulting Lagrangian can be expressed as

$$L(\mathbf{x},\Upsilon,\mu,u) = \sum_{\omega\in\Omega} L_{\omega}(x^{\omega},\Upsilon^{\omega},\mu^{\omega},u^{\omega}) \qquad (20)$$

with $L_{\omega}(x^{\omega},\Upsilon^{\omega},\mu^{\omega},u^{\omega})$ defined by

$$\begin{aligned} L_{\omega}(x^{\omega},\Upsilon^{\omega},\mu^{\omega},u^{\omega}) = & p^{\omega} F(x_{\omega},\Upsilon_{\omega},\mu_{\omega},u^{\omega}) && (21)\\ & + \sum_{i}\left(\sum_{\omega'\in\Omega} p^{\omega'}\Upsilon_i^{\omega}\phi_i^{\omega'} - \Upsilon_i^{\omega}\phi_i^{\omega}\right) && (22)\\ & + \sum_{t\leq 5}\left(\sum_{\omega'\in\Omega} p^{\omega'}\mu_t^{\omega}\tau_t^{\omega'} - \mu_t^{\omega}\tau_t^{\omega}\right) && (23) \end{aligned}$$

$$+\sum_{t\leq 5}\left(\sum_{\omega'\in\Omega} p^{\omega'} u_t^{\omega}\gamma_t^{\omega'} - u_t^{\omega}\gamma_t^{\omega}\right) \tag{24}$$

The corresponding Lagrangian dual problem for problem (5)-(9) is then

$$\min\nolimits_{\phi,\tau,\gamma}\{\mathscr{Z}(\phi,\tau,\gamma) = \max\{\textstyle\sum_{\omega\in\Omega} L_\omega(x_\omega,\Upsilon_\omega,\mu_\omega,u^\omega) : (6)\} \tag{25}$$

Problem (25) 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 $\mathscr{Z}(\phi,\tau,\gamma)$ is required to obtain a subgradient. We note that $\mathscr{Z}(\phi,\tau,\gamma)$ is separable, and reduces to solving $|\Omega|$ problems of manageable size, each of which corresponds to a single scenario. Components of the subgradient vector are then given by $\sum_{\omega'\in\Omega} p^{\omega'}\Upsilon_i^{\omega} - \Upsilon_i^{\omega}$, $\sum_{\omega'\in\Omega} p^{\omega'}\mu_t^{\omega'} - \mu_t^{\omega}$, and $\sum_{\omega'\in\Omega} p^{\omega'}u_t^{\omega'} - u_t^{\omega}$ where Υ_i^{ω}, μ_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 follow Solak et al. (2010), and use a weighted combination of the subgradients from previous iterations in updating the dual variables, such that:

$$\hat{\Gamma}^j = \delta_0\Gamma^j + \delta_1\Gamma^{j-1} + \delta_2\Gamma^{j-2} \tag{26}$$

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:

$$\phi^{j+1} = \phi^j - \frac{\kappa(\bar{L}^j - \underline{L}^j)}{||\Gamma^j||}\hat{\Gamma}^j \tag{27}$$

$$\tau^{j+1} = \tau^j - \frac{\kappa(\bar{L}^j - \underline{L}^j)}{||\Gamma^j||}\}\hat{\Gamma}^j \tag{28}$$

$$\gamma^{j+1} = \tau^j - \frac{\kappa(\bar{L}^j - \underline{L}^j)}{||\Gamma^j||}\}\hat{\Gamma}^j \tag{29}$$

where ϕ and κ, $\kappa < 2$, are constants that can be modified during the algorithm. The values to be used for these parameters 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 still is not fast 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 was possible by establishing the following result about the structure of the optimal investment decisions for the given piecewise linear returns functions.

PROPOSITION 2. *For the extended stochastic DICE model, if* $\lambda_i^{k\omega,*}$ *represent the optimal values for variables* $\lambda_i^{k\omega}$*, then* $\lambda_i^{k\omega,*} \in \{0,1\}$ *for all* k*,* 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: We establish this result through a marginal analysis. First consider technology category 1. Assume that for some k and ω, $0 < \lambda_1^{k\omega,*} < 1$, i.e. the optimal investment is not a vertex value. If the returns function is increasing (decreasing) between the kth and $k+1$st vertex, we can show that it is possible to increase social utility by increasing (decreasing) the investment level to the value at vertex $k+1$ (k).

Suppose we fix all decision variables at their optimal values and increase the investment by Δ units, corresponding to an increase of Δ^{α_1} in the value of the variable α_i^{ω}. This implies a change Δ^y in the production function value y_t^{ω}, where

$$\Delta^y = \frac{(0.8c(\mu_t) - 0.06c(0.5)\mu_t)\Delta^{\alpha_1}}{D(\tau_t)} y_t^g \tag{30}$$

For the given data, it can be shown for all possible parameter and variable values that $\Delta^y \geq \kappa\frac{\Delta}{5}$. Given that maximization of the utility in each period implies the maximization of the net output y_t for that period, it follows that the optimal investment level for technology category 1 must be a vertex value. The same conclusion can be reached by a similar analysis of the case for technology category 2. $\square$

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. This procedure is very effective, and improves the overall solution time significantly as the subgradient iterations will only be implemented over the variables Υ and μ for given investment levels. We implemented this solution methodology to obtain the results described below.

4. Experiments

In this section we discuss the experiments we ran with our stochastic optimization model, where we have based our experimental framework on Nordhaus (2008). In the first subsection we describe a number of alternative policy environments, in Section 4.2 we describe the different risk cases we consider, and in Section 4.3 we briefly discuss assumptions about the opportunity cost of investment.

4.1. Alternative Policy environments

We considered five different policy environments, which we have named DICE optimal, Stern, Gore, Kyoto Strong, and 2 degrees, as well as a baseline no-controls case. Here we describe these policies. For each policy environment we assume there is no knowledge of technological success and damages until year 2055; 2005 abatement is fixed at 0.005; and the planning horizon is 400 years. In DICE optimal there are no other limitations – the model chooses the optimal R&D investment and abatement path. The baseline no-controls case models the levels and growth of major economic and environmental variables as they would occur with no climate-change policies.

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 being the very low discount rate in the former. Thus, this policy is implemented by first running the DICE model at a very low discount rate. This is done by assuming a value of 1 for the elasticity of marginal utility of consumption and a value of 0.001 for the initial pure rate of time preference. We then take the resulting abatement levels and fix those in the model with the default discount rate. For our implementation, we have run two versions of this policy. In one case (we will call Stern fixed), we fix the R&D investment from the run with the very low discount rate. In the other case (we will simply call Stern), we choose the R&D investment in the second run with the DICE discount rate.

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. 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 optimistic, but potentially achievable, international agreement on climate change. This is intended to follow the spirit of Kyoto, but continue on indefinitely and have more and more nations join on through time. In this policy in the original DICE model, abatement is fixed for the first 150 years. To be able to model the learning about climate damages and R&D, we altered this, 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 2 Degrees policy simply adds a single constraint that limits the temperature increase to 2°C. In the DICE model 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 of the DICE model. 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.

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

Table 7 Probability distributions defining climate change damage uncertainty.

4.2. 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 7 shows the five cases we consider.

Each probability distribution is given a name in the top row. The second row shows the probabilities of each outcome in that distribution. The third row shows the percent GDP loss given a 2°C increase in global mean temperature as calculated using equation (1) in the DICE model above. 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 last row shows the value of the parameter π for each respective outcome. In previous work we have worked with an MPS around this value. However, the damages in the DICE model are concave in π (see equation (1) above), therefore the expected GDP loss decrease is an MPS around π.

Moving from case 1 to case 4, each distribution is an MPS of the distribution to the left. Case 5 is intermediate in risk between case 1 and 3, but cannot be compared directly with case 2. Case 5 is based on the probability of catastrophic damages reported in Nordhaus (1994).

4.3. Opportunity Cost

The R&D funding levels used in the elicitations and reported in Tables 1-3 represent the amount of money going into the hands of high quality researchers in the appropriate areas (Baker et al. 2008, 2009a,b). 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

Investments ($ million)									Total Inv
Pre C	Chem L	Post C	LWR	HTR	FR	Org.	Inorg.	3rd g	($bil)
386	56	519	346	3089	15443	830	77	386	21.132
386	56	519	346	3089	0	116	77	386	4.975
154	56	519	346	3089	0	116	77	386	4.743
39	56	224	346	0	0	0	77	0	0.742

Table 8 Allocation of total investment under different optimal investment values

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. $\kappa = 4$). 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 in equation (2), and discuss it as part of our analysis in Section 5.

5. Results

In Sections 5.1 – 5.4 we present our results.

5.1. Optimal Investment in R&D

We find that the optimal investment in R&D is quite robust, both across different policy environments and across different risk cases. Fig 5 illustrates the optimal investment across risk and policy enviornments. The horizontal axis represents the amount of high damages in each risk case, considering the no-, medium-, and high-risk cases. We see that the most common optimal investment is equal to $4,975 million. In Table 8 we show the allocation of the total investment in research areas for the optimal levels of investment in different cases. Hence, an investment of $4,975 million represents a high investment in all CCS technologies, as well as LWR, HTR, inorganic and 3rd generation solar, plus low investment in organic solar technology. This is the optimal investment for Gore and 2 degrees at all risk levels, plus for DICE-Optimal, Stern, and Kyoto Strong for risk cases 1 and 2. For the last three, the optimal investment reduces slightly under high risk (case 3), to $4,743. The intermediate risk case also has this investment level for DICE-Optimal.

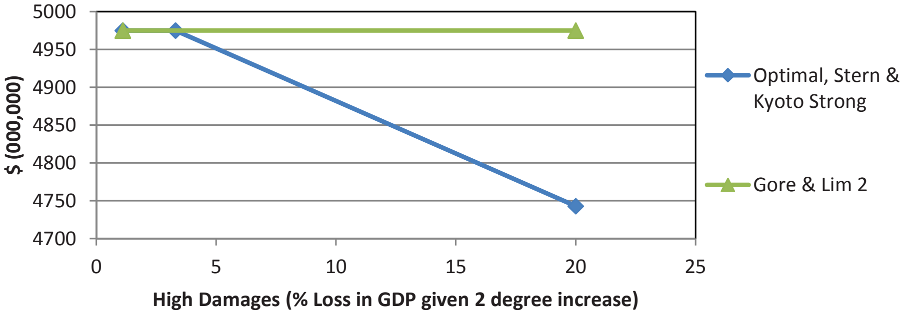

Figure 5 **Optimal investment across risk and policy scenarios**

Not shown is the DICE-optimal investment under very high risk. This falls to $742 million, representing high investments in LWR, inorganic solar and chemical loop, as well as medium investment in post-combustion and low investment in pre combustion technologies. 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 follows the reasoning given above: When damages are very high, abatement is at 100% with or without technical change. Technical change does serve to reduce the cost of full abatement, but this only happens with a low probability, thus the expected value of technical change is lower under very risky cases, and optimal investment is decreased. Thus, it appears that technical change does not provide much of a hedge against riskiness in damages.

The optimal investment is similarly robust to assumptions about the opportunity cost of R&D investments. We find that if the opportunity cost is between $1-4$, then the optimal investment is stable as above. If it is between $6-10$, then the optimal investment is slightly lower at $4,743 million. Thus, the investment is not very sensitive to assumptions about opportunity cost.

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.

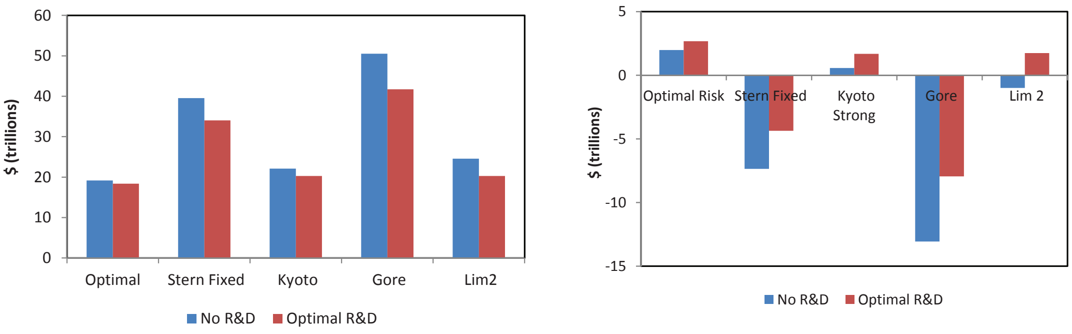

(a) Effect of R&D on total costs

(b) Expected utility of policy interventions

Figure 6 **Costs and utilities of policy interventions**

5.2. Impact of R&D on Expected Policy Costs

Figures 6(a) and 6(b) show the impact of R&D on the different policy environments. The left panel shows the total cost of each policy with and without R&D; the right panel shows the expected utility of each policy intervention with and without R&D.

The vertical axis in Figure 6(a) represents the net present value of the cost of abatement plus the cost of climate damages in trillions of 2005 dollars. The vertical axis in Figure 6(b) represents expected utility translated into trillions of 2005 dollars.[1] Each bar represents the extra utility gained (or lost) from the baseline by implementing the policy intervention. The lighter 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 darker 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 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. Our focus is not on the evaluation of the policies themselves, but rather on the role of R&D in mitigating the costs of the environments.

The first result here is that the availability of R&D is more valuable in 2nd best policy environments. The absolute value of having R&D is greater in the non-optimal environments and is

---

[1] Each bar represents the difference in expected utility between the named policy and the baseline no controls case. We normalized the values by multiplying this by the ratio of the cost difference between the no-R&D no risk DICE optimal case and the no controls case, divided by the difference in expected utilities of the same two cases.

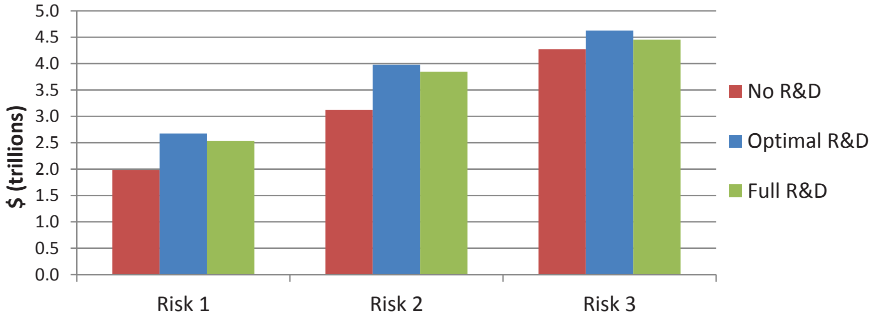

Figure 7 **Expected utility of policy intervention for the DICE optimal policy**

greatest in the Gore environments. Of particular interest are the Kyoto Strong and 2 degree 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 2 degree limit goes from being a net loss to a net benefit with R&D.

Figure 7 shows the expected utility of policy intervention for the DICE-optimal policy, comparing no R&D, optimal R&D, and full R&D, 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 impact.

5.3. Impact of R&D on the range of scenarios

Figure 8 shows the range of emissions paths in the risk 1 case for DICE optimal when R&D is invested optimally.

There are a total of 28 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 emissions between 70-75% in 2105 is 43%. The lowest line is the case where all technology fails. In the rest of the figures we will just show the range of paths without associated probabilities.

Figure 9 compares the range of emissions paths from the DICE optimal, the Stern fixed, and the Kyoto Strong policies. The high group are the paths from Stern, the lower group from DICE

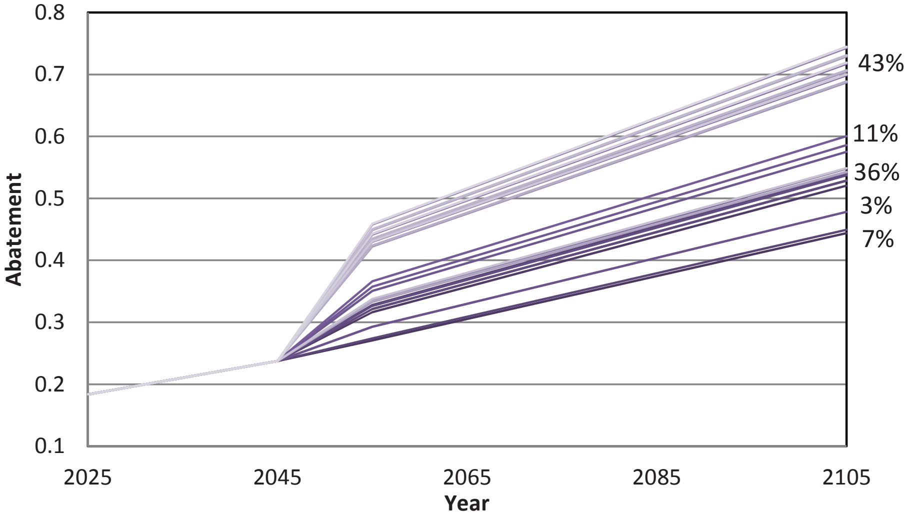

Figure 8 **Range of emissions paths in the risk 1 case for DICE optimal**

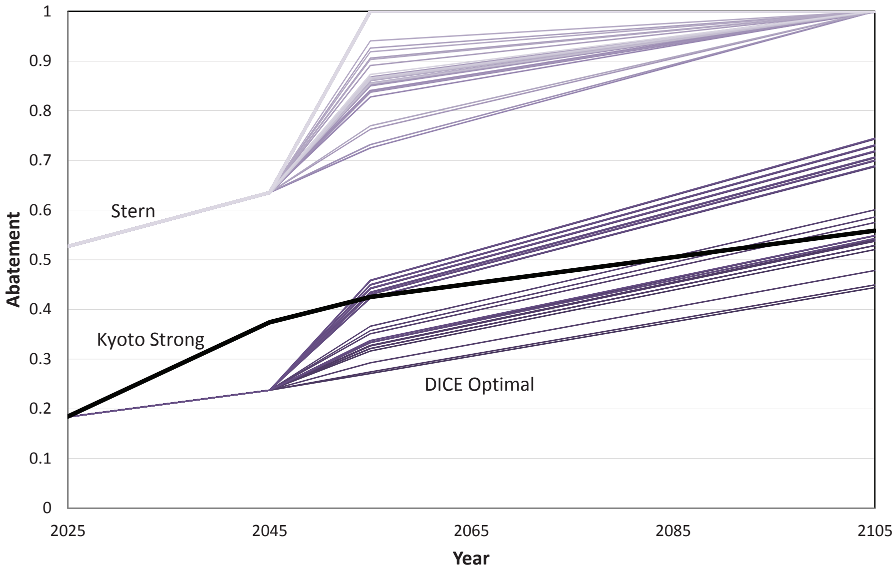

Figure 9 **Comparison of the range of emissions paths from DICE optimal, Stern fixed, and Kyoto Strong**

optimal, and the thick line is for Kyoto Strong.

Note that in the absence of technical change the Kyoto Strong path is clearly higher than the optimal path; with technology, however, it falls about in the middle of the optimal emissions paths. Thus, the presence of R&D greatly enhances the fixed emissions path prescribed by Kyoto Strong.

Figure 10 shows the range of temperature paths in the risk 1 case for DICE optimal (the higher

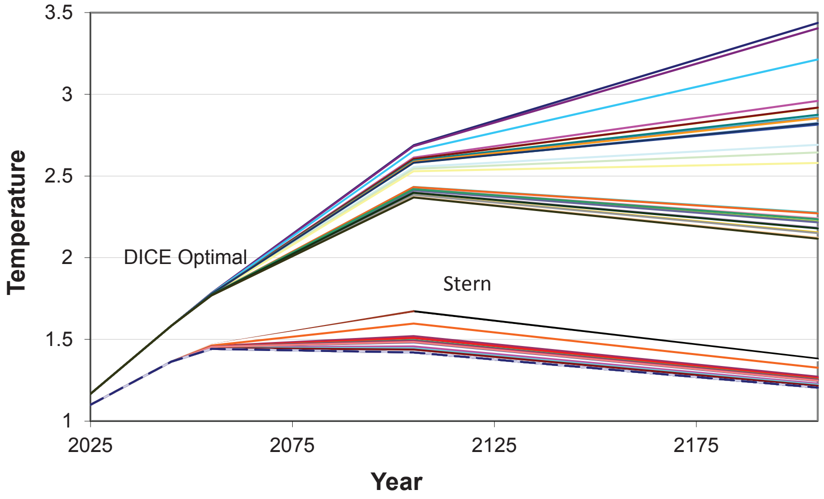

Figure 10 **Range of temperature paths in the risk 1 case for DICE optimal and for Stern**

group) and for Stern (the lower group). From this figure we can make three observations. First, the DICE optimal path is always above 2 degrees while the Stern path is 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 find this goal to be infeasible to solve, 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 11 shows the abatement costs for all scenarios for the Risk 1 case of DICE optimal and Stern. The figure on the left shows the years 2050 to 2105 at a higher scale than the figure on the

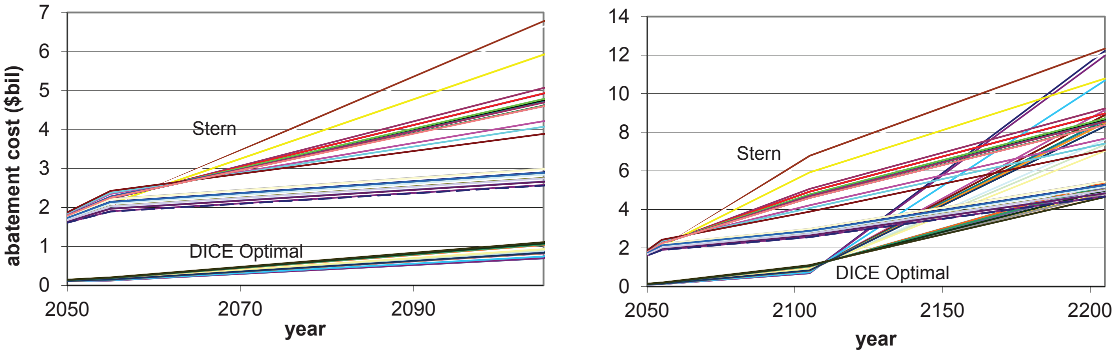

Figure 11 **Abatement costs for all scenarios for the Risk 1 case of DICE optimal and Stern**

right, which goes out until 2205. We see that until 2105 R&D has a much larger impact on Stern than on DICE optimal. In fact, the figure on the left is almost the opposite of Figure 10 above 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.

5.4. 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 12 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 given an optimal investment in R&D is less than $170 trillion 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 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 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.

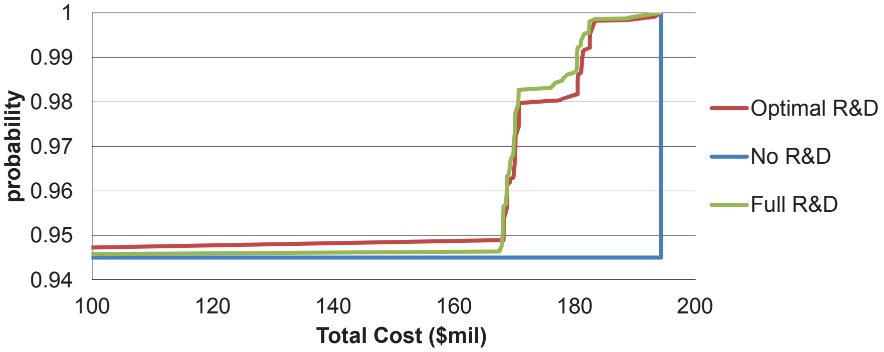

Figure 12 **Cumulative distribution function comparing different R&D policies**

There is about a 3.5% chance that damages will be below $170 trillion. Thus, R&D provides risk reduction.

There is very little difference between optimal R&D and full R&D. There is some cost to over-investing in R&D in the full investment case, thus the optimal R&D line is higher on the left side of the graph. On the other hand, full R&D provides some risk reduction in the high risk case, and thus lies above the optimal R&D line at some places.

Figure 13 shows a CDF for the Risk 3 case, for DICE optimal and for Stern. On the left side of the graph, we see the cost of the Stern policy – if damages turn out to be very low, Stern will cost about $19.9 trillion versus the DICE optimal cost of $8.4 trillion. On the right side of the graph we see how the policies compare when damages turn out to be high. In this case, the Stern policy is better, even when evaluated within the DICE model. Note that in the Stern case, when damages are high, there is a probability of limiting the total damage to $125 trillion; whereas the DICE optimal case will have total costs of at least $169.9 trillion. This figure shows that there is no first- or second-order stochastic dominance between the two policies. A very risk averse decision maker might choose the Stern policy over the DICE optimal policy. That is, the DICE optimal policy is optimal for a specific utility function. But, for decision makers that are very risk averse, even when evaluating the Stern policy at the lower DICE discount rate, it may be optimal to choose Stern. Thus, Stern can be looked at as a response to risk aversion.

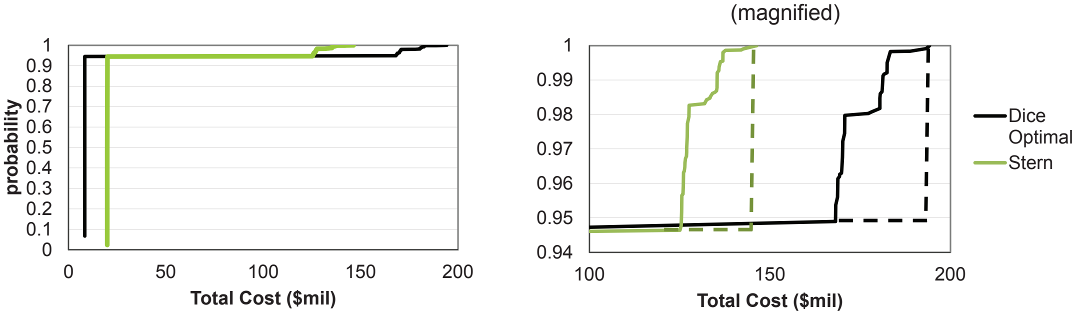

Figure 13 **Cumulative distribution functions for Risk 3 DICE optimal and Stern cases**

The right panel of Figure 13 magnifies the high end of the horizontal axis. The dotted lines show what the CDFs would look like in the absence of R&D. R&D reduces total costs in the high damage case, the very case where society cares the most, and so R&D can be looked at as a risk reduction response.

6. Conclusions and Policy Insights

We have implemented data-based uncertainty on the returns from R&D into a stochastic version of an IAM in order to get insights about the optimal R&D portfolio in the face of climate change. We built on previous work in which we used a detailed project-by-project model to determine the optimal mix of projects at a set of given budget levels. This allowed us to develop a simplification of the R&D description to implement in a much more detailed and dynamic model.

Our first set of conclusions relate to the R&D investment itself. We found that focused investments in Nuclear LWR, HTR as well as in CCS and solar technologies are very robust given our data. One result of using *data* (rather than theoretical explorations), is that 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. There are two important caveats on these particular recommendations that need to be taken account in a final investment decision. First, 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. Second, we, as a society, have only a weak understanding of how intermittent renewables, such as solar PV, 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 PV will be larger than the current models show, especially if simultaneous investments are made in the grid and grid integration. Thus, these two aspects should be considered in a final portfolio allocation.

We have shown that a larger-than-optimal investment in technology is less costly than a smaller-than-optimal investment; and a smaller-than-optimal investment has larger implications if damages turn out to be higher than expected. Thus, it appears that policy makers should prefer to err on the high side rather than the low side of R&D investment.

The optimal investment in R&D appears to be surprisingly robust to the riskiness of climate damages. We do see a lower investment when there is a very small probability of very high damages. But, given the level of robustness, the result above, and the deep uncertainty about climate damages (i.e. that we really don't know the probability distribution), these things lead to 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 Nordhaus' (since very few people 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, surprisingly, have little impact on the results, and therefore are not of great concern.

Our second set of conclusions relate to the role of R&D and uncertainty in the different policy environments that we consider. First, we see that R&D has more value in "2nd best" policy environments. The difference between R&D and no-R&D was lowest in the DICE-optimal policy. This is important since, in the real world, given political reality and vast uncertainties, we will surely be implementing a 2nd-best policy. Thus, the value of government R&D will be higher than most current models will show. Of particular note is the role of R&D in the Kyoto-Strong policy, since this is an attempt to model a strong but potentially implementable policy. We see that without R&D this policy is barely better than doing nothing; with R&D it becomes significantly beneficial.

While the optimal investment in R&D is fairly robust across policy environments, we see that the *role* of R&D varies considerably, both with the policy environment and with the damage-risk case. In policy environments in which emissions are fixed or tend to be very high (near 1), R&D primarily has a cost-side benefit: the environmental variables are less effected while the cost of abatement is significantly effected. 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 effected, while the cost of abatement has only small effects, and in fact sometimes is higher given a much higher level of abatement.

Finally, we compared the full risk profile of the DICE-optimal and Stern policies, with and without R&D, and found that none of the policies shown second order stochastically dominate another policy. In general, more risk averse decision makers would lean toward higher levels of R&D and more stringent policies such as Stern. Thus, these policies can be seen as a response to risk aversion.

Overall, our analysis is a comprehensive and integrated approach to a highly stochastic problem of R&D portfolio analysis for energy technology investments in response to climate change. The fact that our analysis elicits and uses actual data, and then implements it in a highly complex stochastic dynamic assessment model in a tractable way is a significant contribution. Hence, the performed experiments and the resulting policy implications are of high importance for policy making.

Acknowledgments

This research was partially supported by NSF under award number SES-0745161, and the Precourt Energy Efficiency Center at Stanford University. We thank Yiming Peng and Olaitan Olaleye for excellent research assistance, and Haewon Chon for providing some updated MiniCAM results.

References

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

Baker, Erin, Haewon Chon, Jeffrey Keisler. 2008. NuclearPower: Combining expert elicitations with economic analysis to inform climate policy. Available at SSRN: http://ssrn.com/abstract=1407048.

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

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

Baker, Erin, Ekundayo Shittu. 2008. Uncertainty and endogenous technical change in climate policy models. *Energy Economics* **30**.

Baker, Erin, Senay Solak. 2010. Climate change and optimal energy technology R&D policy. *European Journal of Operations Research* **under revision**.

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.

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(Supplement 2)** s64 – s81.

Clarke, L., P. Kyle, M.A. Wise, K. Calvin, J.A. Edmonds, S.H. Kim, M. Placet, S. Smith. 2008. CO2 emissions mitigation and technological advance: An updated analysis of advanced technology scenarios. Tech. Rep. PNNL-18075, Pacific Northwest National Laboratory.

Clarke, Leon E., James A. Edmonds, Henry D. Jacoby, Hugh M. Pitcher, John M. Reilly, Richard G. Richels. 2007. Scenarios of greenhouse gas emissions and atmospheric concentrations. Tech. Rep. Synthesis and Assessment Product 2.1a, U.S. Climate Change Science Program.

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* .

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

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.

Nordhaus, W. 1993. Rolling the DICE: An optimal transition path for controlling greenhouse gases. *Resource and Energy Economics* **15** 27–50.

Nordhaus, William. 2008. *A Question of Balance: Economic Modeling of Global Warming*. Yale University Press, Connecticut, USA.

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

Nordhaus, William D. 2002. Modeling induced innovation in climate change policy. Arnulf Grubler, Nebojsa Nakicenovic, William D. Nordhaus, eds., *Technological Change and the Environment*. RFF and IIASA, Washington D.C. and Laxenburg, Austria, 182–209.

NRC. 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.

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

Rothschild, Michael, Joseph Stiglitz. 1970. Increasing risk I: A definition. *Journal of Economic Theory* **2** 225–243.

Solak, Senay, Erin Baker. 2010. Convexity analysis for the dice model. *working paper* .

Solak, Senay, John-Paul B. Clarke, Ellis L. Johnson, Earl R. Barnes. 2010. Optimization of r&d project portfolios under endogenous uncertainty. *European Journal of Operational Research* **207**(1) 420 – 433.

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