Climate Change and Optimal Energy Technology R&D Policy: Implementing Uncertainty and Learning in Policy Analysis

Erin Baker* and Senay Solak†

DRAFT prepared for University of Arizona Economics Seminar

Abstract

Global climate change presents a classic problem of decision making under uncertainty with learning. Much work has shown that explicitly accounting for uncertainty and learning in climate change can have a large impact on optimal policy, especially technology policy. We provide a framework for collecting and implementing probabilistic data into policy analysis. We consider the optimal energy technology R&D portfolio in the face of cimate change, among three key power technologies: solar PV, carbon capture and storage, and nuclear power. We find that the composition of the optimal portfolio (for a given budget level) is quite robust; but the optimal level of investment depends on uncertainty. In general, the *a priori* value of technical change appears very different when we recognize that later stage decisions are available.

JEL classification: D81;O32; Q54; Q55; Q58

Keywords: Climate change; Information and uncertainty; Environmental policy; R&D

*Correspondence Address: Erin Baker, 220 Elab, University of Massachusetts, Amherst, MA 01003; edbaker@ecs.umass.edu; 413-545-0670

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

1 Introduction

While scientists largely agree that humans are changing the climate, there is a great deal of uncertainty about the degree to which emissions of the greenhouse gasses that cause global warming will cause damages in the future. This is part of what is causing a very large debate about how best to proceed in the fight against global climate change. Possible near term policy responses include both restrictions on emissions (through emissions limits or taxes) and investment in environmentally friendly technologies. A number of researchers have investigated the question of how the presence of uncertainty and learning impacts near term optimal policy. The answer to this question seems to be "it depends": optimal near term decision variables, such as abatement or R&D investment, may increase or decrease with increases in risk or increases in learning. Thus, the next step is to try to characterize the uncertainty that we are facing and implement this into policy models.

We focus on socially optimal near term investments in energy technology R&D, with an emphasis on the composition of the optimal portfolio. Determining this allocation of investments involves a number of issues, specifically – based on the research mentioned above – it requires explicitly incorporating uncertainty over multiple dimensions [**?**]. The process of R&D is inherently uncertain – we cannot predict whether any particular program will be successful, or the degree to which it will meet or exceed goals. In the case of climate change, we also have deep uncertainty on the benefits side – there is considerable uncertainty about the damages that will be caused by climate change, and hence, the benefits from reducing emissions. This feeds back, to create uncertainty about the value of having any particular technology available. In this paper we perform a portfolio analysis to get insights about the optimal R&D portfolio, and how it changes with increasing risk in climate damages.

In Section 2 we present a simple model of decision making under uncertainty in climate change and review the literature in the context of this model. In the rest of the paper we present a framework for collecting and implementing probabilistic information on technical change when there is uncertainty and learning. In Section 3 we describe our previous work, combining probabilistic data from expert elicitations with economic modeling to determine probabilities over Marginal Abatement Cost Curves. We focus on three key climate change energy technologies – solar photovoltaics (PV), nuclear power, and carbon capture and storage (CCS). In Section 4, we describe our portfolio model, and in Section 5 we present results that indicate (1) the value of moving beyond theoretical analysis to data-based analysis; and (2) the importance of explicitly modeling uncertainty and learning. In particular, in explicitly modeling a second stage abatement decision, we capture an environmental-side benefit that is missed in many analyses: when technical change is successful it can lead to optimally higher abatement. Recognizing this leads to an overall higher value to R&D. We conclude in Section 6.

2 The Ambiguous Impacts of Uncertainty and Learning in Climate Change

[1]A number of papers have addressed the role of uncertainty and learning in climate change, with a focus on determining what impact uncertainty and learning have on near term decisions. The overall result is that the impact of uncertainty and learning is ambiguous, as mentioned in the introduction. This is true for both optimal abatement and optimal investment in R&D. In this section we discuss why uncertainty and learning leads to ambiguity and review the literature on the subject. We start by defining what we mean by increasing uncertainty or risk, and then go on to use a simple model to illustrate the general insights.

2.0.1 Definition of Risk

We focus on increasing risk or mean-preserving-spreads (MPS), as defined by Rothschild & Stiglitz [**?**]. Let Z represent a random variable.

Definition 1 *Z is **riskier** (or **more uncertain** or **more variable** or an **MPS**) than Z' iff $\mathrm{E}_Z\, U(Z) \leq \mathrm{E}_{Z'}\, U(Z')$ for all concave U.*

By definition then, risk is exactly what risk-averters don't like.[2] We focus on MPS because they do not conflate the impact of an increase in mean with an increase in risk; and they are relatively easy to characterize and interpret. There is also an added benefit to considering MPS – there is a strong parallel between MPS and an increase in informativeness, in the Blackwell [**?**][**?**] sense (See [**?**]). Thus, through studying what kinds of MPS increase the optimal value of a decision variable, we can infer what kinds of signals increase the optimal value of a decision variable. Rothschild and Stiglitz [**?**] showed that this definition is equivalent to Z having more weight in the tails than Z' (characterized by what are often called the integral conditions).

2.1 An illustration

Mathematically, it is the shape of the marginals that drive the ambiguity. Here we illustrate the basic idea through a very simple model of climate change decision making with uncertainty in the damages. For this model we ignore uncertainty in technical change and represent uncertainty in climate change through a single random variable Z that impacts the damage curve. We concentrate on two first period decisions: abatement, μ, and R&D investment, α. The second period decision that provides the flexibility is second period abatement, μ_2.

$$\min_{\mu \leq 1, \alpha} c_1(\mu) + g(\alpha) + \mathrm{E}_Z \left[\min_{\mu_2 \leq 1} c_2(\mu_2; \alpha) + D(S - \mu - \mu_2; Z) \right] \tag{1}$$

[1]Much of this section is based on [**?**]

[2]This definition differs from what is often called Second Order Stochastic Dominance, in that SOSD is defined only for *increasing*, concave U. The result of this different definition is that "increasing risk" only orders random variables with equal means; SOSD orders a larger set. In general, the more restrictions that are put on the set of functions U, the larger is the set of probability functions that can be ordered; and vice-versa. See Athey [**?**] and Osborn [**?**] for a discussion of this relationship.

where μ, μ_2 are first and second period abatement (measured as the fraction of emissions reduced below the business-as-usual level of emissions); α is the impact on the second period abatement cost curve from technological change; S is the current stock of emissions; Z is the (positive) random variable that impacts damages; c_1, c_2 are the abatement cost functions in the first and second period; g is the cost of achieving technical change equal to α, and D is the damage from global climate change. We make the standard assumptions that g, c_1, c_2 are increasing and convex; and that D is increasing and convex in the stock $s = S - \mu - \mu_2$. For this simple model we assume that technical change takes a specific form, namely that it will decrease the cost of second period abatement proportionally: $c_2(\mu_2; \alpha) = (1-\alpha)c_2(\mu_2)$.[3]

Abatement in the second stage is assumed to be optimal, and to depend on climate damages. Thus, this model represents uncertainty with perfect learning. Formally, the optimal interior value of $\mu_2^*(z)$ (where lower-case z represents a realization of the random variable Z) satisfies the first order condition

$$\frac{\partial c_2}{\partial \mu_2} = \frac{\partial D}{\partial s} \tag{2}$$

We may, however, have a corner point solution, where $\mu_2^*(z) = 1$ and $\frac{\partial c_2}{\partial \mu_2} \leq \frac{\partial D}{\partial s}$. We define $\bar{z}$ as the "full abatement point", the level of damage that induces the corner point solution: $\mu_2^*(z) = 1$ for all $z \geq \bar{z}$.

The first order conditions for μ and α are as follows:

$$c_1'(\mu) = \mathrm{E}_z\left[\frac{\partial D(S - \mu - \mu_2^*; Z)}{\partial S}\right] \tag{3}$$

$$g'(\alpha) = \mathrm{E}_z\left[c_2(\mu_2^*)\right] \tag{4}$$

where μ_2^* is second period optimal abatement.

The right-hand-sides of (3) and (4) are the marginal benefits to near term abatement and to R&D. The right-hand side of (3) is the marginal change in damages due to a reduction in the stock of emissions, assuming that second period abatement is optimal; we will call this optimal marginal damages from now on (to distinguish from marginal damages, holding second period abatement constant).[4] The marginal benefits to abating in the first period are equal to the expected optimal marginal damages. On the other hand, the marginal benefits to investing in R&D are equal to the expected cost of optimal second period abatement.

Any change in the probability distribution of Z that increases (decreases) the expected marginal benefits will cause the optimal value of the decision variable to increase (decrease).[5] Expected marginal benefits will increase (decrease) for all increases in risk if and only if the marginal benefits are convex (concave) in Z.[6] This is a direct result of the definition of increasing risk. Thus, we need to investigate the shape of optimal marginal damages and the shape of the cost of optimal second period abatement, as a function of the random variable Z.[7]

---

[3] See[?][?][?][?][?][?][?][?] for papers that make similar assumptions about technical change.

[4] Formally, marginal damages are defined as $\frac{\partial D(s,z)}{\partial s}$; optimal marginal damages are defined as $\frac{\partial D(S-\mu-\mu_2^*,z)}{\partial S}$.

[5] Throughout the paper we will use the term "increasing" to mean non-decreasing, and will say "strictly increasing" when that is what we mean.

[6] To be precise, define $x(F)$ as the optimal value of the decision variable x given the probability distribution

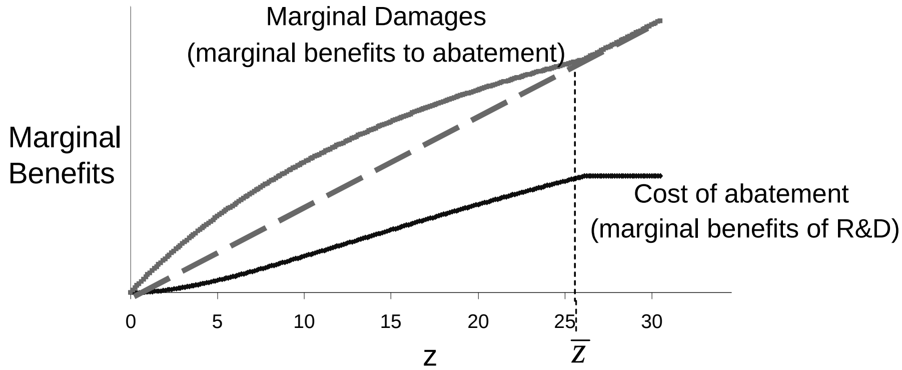

Figure 1: The marginal benefits of abatement and R&D as a function of multiplicative uncertainty on damages, Z. The marginal benefits to abatement are equal to $ZD'(S-\mu-\mu_2^*)$; the marginal benefits to R&D are equal to $c(\mu_2^*)$ .The dashed line represents the marginal benefits to abatement, holding second period abatement constant at full abatement $\mu_2 = 1$: $ZD'(S-\mu-1)$.

Figure 1 illustrates the shape of the marginal benefits to abatement and to R&D, under the assumptions that costs and damages are quadratic, damages are linear in z, and abatement in both periods is restricted to be less than or equal to 1. The marginal benefits to abatement (the upper, solid line) are equal to $zD'(S-\mu-\mu_2^*)$ and here are graphed as a function of z. Note that changes in z have a direct effect on optimal marginal damages (increasing linearly in this case), and an indirect effect through optimal second period abatement μ_2^*. They are concave in z, until the full abatement point, $\bar{z}$. This reflects the fact that before full abatement, second period optimal emissions decrease as damages get worse. Thus, optimal marginal damages are concave (in fact this is true as long as $D''' \leq 0$). When the full abatement point is reached, however, second period emissions can no longer be reduced and optimal marginal damages become linear in the variable z. The slope of the linear part of the curve is always steeper than the slope of the concave part of the curve, thus the curve is reverse S-shaped with inflection at $\bar{z}$. (See Baker, 2005 for proofs). The key points here are 1) marginal damages are initially concave since there is room to respond optimally in the second period; 2) there is an inflection point at full abatement; and 3) the shape of the marginal benefits to abatement after the inflection point depends only on the direct effect of the random variable on the marginal damages.

The marginal benefits to R&D are equal to $c_2(\mu_2^*)$. This is impacted indirectly by z, since z impacts optimal emissions, μ_2^*. This marginal has a shape that is in some sense opposite to the marginal for abatement. The cost of optimal second period abatement is convex at $z = 0$, since $c_2(\cdot)$ is convex and μ is increasing in z. However, optimal second period abatement is concave in z – it becomes increasingly expensive to reduce the next unit of emissions, therefore the marginal reduction in emissions slows down. This leads to an inflection point $\hat{z} \leq \bar{z}$ after

F. Then we say x is **increasing (decreasing) in risk** if $x(G) \geq (\leq)\, x(F)$ whenever G is riskier than F.

[7] Note that the impact of risk on the net marginal benefits (the benefits minus the costs) are exactly equal to the impact of risk on the gross marginal benefits.

which the cost of optimal second period abatement is concave[8]. Finally, the cost of optimal second period abatement is constant for $z \geq \bar{z}$ since second period abatement is constant for that range. This leads the marginal benefits of R&D to be S-shaped. The key points here are 1) optimal cost is initially convex if cost is convex and second period abatement is increasing rapidly near the origin; 2) optimal cost will generally become concave in high damages, since optimal second period abatement will be strongly concave (as it approaches a maximum level); and 3) if a full abatement point exists, then optimal cost will be constant after this point.

The figure shows that for both control variables the marginal benefits are neither convex nor concave under even these extremely simple assumptions, and thus both optimal first period abatement and optimal R&D will increase with some increases in risk and decrease with other increases in risk.

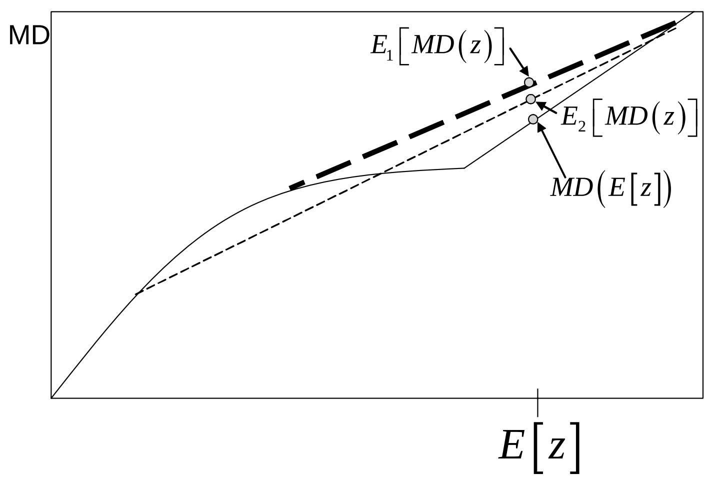

Figure 2: The thick dashed line represents an MPS in the random variable Z; the thinner dashed line represents an additional MPS. Expected marginal damages first increase, and then decrease with the successive MPS.

A prominent example of illustrating the ambiguity of optimal abatement is found in Ulph and Ulph [?]. They go on to show in Lemma 3iii that if the mean damages are high, $\mathrm{E}[z] > \bar{z}$, then abatement will be higher under risk than under no risk.[9] To understand how this result is related to the shape of the marginal damages see Figure 2. Note that any simple MPS around any point above $\bar{z}$ will lead to (weakly) higher expected marginal damages due to the convexity induced by the full abatement point. This is illustrated in this case by both dotted lines. However, note that the thinner dotted line represents a riskier distribution than the thicker, and yet the expected marginal damages decrease with a move to the riskier distribution. Baker [?] provides a related, but more general result, showing that optimal abatement will increase in MPS that stretch the tail of the distribution.

[8]Specifically, if $c_2(\mu) = a\mu^2 + b\mu$ and $D(S-\mu) = (S-\mu)^2$ then $\bar{z} = \frac{a+\frac{b}{2}}{S-1}$ and $\hat{z} = \min\left[\frac{a\left(aS-\frac{b}{2}\right)^2+ab}{2aS+b}, \bar{z}\right]$

[9]Their result is stated in terms of learning versus no learning. Given their linear model, their result is true if and only if abatement is higher under risk than under no risk. See Baker [?] for details of the equivalence between increases in risk and increases in learning.

Webster [?] analyzes a similar model and concludes again that the impact of risk is ambiguous, and that what matters is how the probability is redistributed. Kolstad [?], Gollier, Julien, and Treich [?], Karp and Zhang [?], and Baker [?] each consider the impact of an increase in learning, and arrive at varying conclusions, finding conditions under which learning causes optimal emissions to increase or decrease.

There is a very recent literature investigating the optimal investment in technology R&D in the face of uncertainty. Some papers consider uncertainty in the climate damages ([?], [?], [?], [?]), some consider uncertainty in technological change ([?], [?], [?]), and one paper considers both [?]. Again, we find that the impact of uncertainty and learning is ambiguous on investment into technical change. However, these papers tend to show that uncertainty in technological change has a quantitatively larger impact on optimal actions than does uncertainty in climate damages, and that the optimal investment in R&D is often much higher when uncertainty is explicitly included. However, all of these studies consider investment in one technology at a time, rather than a portfolio of technologies.

A small number of papers have studied the impact of uncertainty on a *portfolio* of technologies. Two studies that are closely related to this paper are [?] and [?]. They consider the question of the optimal R&D portfolio when there is uncertainty in both technological change and climate damages, with a focus primarily on the drivers of diversification in the portfolio. They show that it is not enough to just consider the potential value of new technologies, but that the uncertain relationship between program funding and effectiveness is just as important.

Finally, [?] considers optimal abatement and optimal R&D investment simultaneously, and finds that abatement and investment in technology R&D may be risk-substitutes in many cases: changes in risk that optimally increase one, decrease the other. Abatement is a hedge against *catastrophic* damages – damages that make society want to abate more than the total flow of emissions. If there is a possibility that damages will turn out to be catastrophic and society may find itself up against a wall, there is value in reducing emissions now in order to preserve the flexibility to respond in the future. On the other hand, if damages turn out to be just a little worse than expected, society can simply abate more in response – there is value to waiting to learn. Technology R&D is a hedge against damages that may be a little worse than average. The reason is that if damages are a little worse than expected, but not catastrophic, then an incremental improvement in technology leads to two benefits: the cost of abating a given amount is lower *and* the optimal abatement is higher. The double payoff is what makes the investment more attractive. However, in the event of a catastrophe, improvement in technology will not lead to higher abatement, it will only lead to cost savings. Since the probability of this cost savings is lower under riskier damages, the optimal investment is smaller.

2.2 Summary

In this section we have shown that the impacts of uncertainty and learning on near term policy are ambiguous, and therefore, theoretical investigations will generally not be able to determine even the direction of the impact. Moreover, it appears that the role of uncertainty and learning may be more important in determining optimal technology policy than it is for optimal abatement policy. A number of papers have performed computational investigations indicating that the optimal investment in R&D may be significantly different when we explicitly include uncertainty and learning in the model. However, to date, the probability distributions over technical success used in the papers have been purely theoretical, intended to test whether uncertainty was important. They have not been intended to specifically inform policy. In the next section we describe a framework that combines economic analysis with decision analysis in order to inform policy using the best available information to derive probability distributions.

3 A Framework for Data-based Policy Analysis under Uncertainty

In this section we describe a framework combining Decision Analysis and Economic Analysis to collect and implement probabilistic data on damages and technical change into climate policy analysis models.

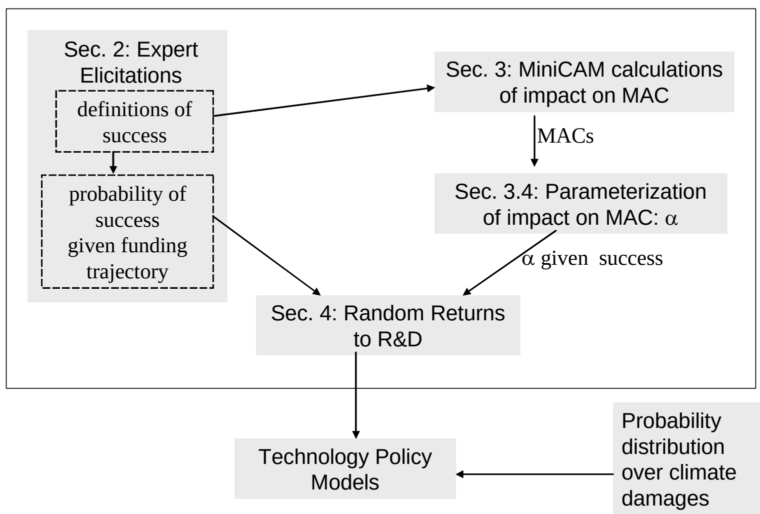

Figure 3: A schematic representation of the flow of data in the framework. The elements inside the box have been discussed in detail in previous work; the elements outside the box are discussed in detail in this paper.

Figure 3 illustrates the flow of the data in this framework; the actions placed within the box have been discussed in detail in [?], [?], and [?]; the actions outside the box are ongoing work that will be discussed in some detail in Section 4 below. In this section we provide on overview of our previous work. We start by discussing the importance of the Marginal Abatement Cost Curve (MAC). We then discuss the process of collecting probabilistic data on technological change through expert elicitations. The products of the elicitations include explicit definitions of success for each technology, and probabilities of success for given funding trajectories. In Subsection 3.3, we translate these definitions of technological success into economically useful parameters. To do this we use a technologically detailed Integrated Assessment Model to determine how the technologies would impact the MAC, if they achieve success as defined. We can then combine the probabilities with the derived MACs to have probability distributions over MACs conditional on different funding scenarios. In Section 4, we implement this data in a stochastic programming model to derive results and insights about the optimal portfolio of technologies.

3.1 Marginal Abatement Cost Curves

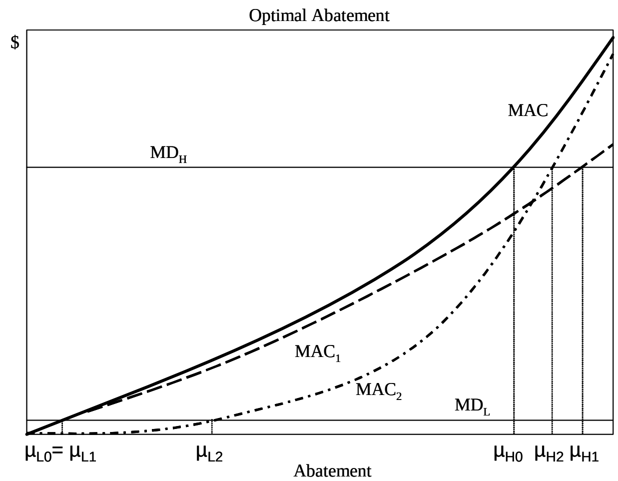

Figure 4: Stylized representations of technical change impact on the MAC; and resulting optimal abatement levels.

The uncertainties in both climate damages and in technical change are dynamic, in that we expect to learn more about each as time goes on. The value of a particular R&D program for a particular technology depends not only on whether the technology development is successful, but may depend on the severity of climate change damages in the future. Some technologies, such as improvements in fossil fuel efficiencies, may have the largest impact if climate change turns out to be mild and only small reductions in emissions are called-for. At very high abatement levels society will tend to substitute away from fossil fuel, and thus improvements in those technologies will have less impact. Other technologies, such as electric vehicles, may have the most impact if climate change turns out to be very severe, calling for an almost total reduction in greenhouse gas emissions. Electric vehicles may not be adopted at low abatement levels, but may prove to be a widespread alternative at very high abatement levels.

The key point is that technical change has two effects: it reduces the cost of abatement for a given level of abatement, but it also changes the optimal amount of abatement. One way to capture both of these effects is to focus on the Marginal Abatement Cost Curve (MAC) [?]. This is the curve that reflects the cost of reducing emissions by an additional ton. Figure 4 illustrates how the impact of technical change on optimal abatement varies by technology and by the severity of marginal damages. The solid upward sloping line represents the original MAC. The two dashed lines represent different types of technical change. The horizontal lines represent two levels of marginal damages. On the horizontal axis we show the optimal level of abatement in each case, where μ_{ij} represents optimal abatement given damages $i = H, L$ and MAC curve $j = 0, 1, 2$. The cost of abatement is represented by the area under the curve to the left of the abatement line. Note that the technical change embodied by MAC_1 has no effect on optimal abatement when marginal damages are low, but a significant effect when damages are high; the impacts of MAC_2 on optimal abatement are nearly the reverse. By paying attention to the impact of technology all along the curve (rather than just a point estimate), we gain information about how optimal behavior will change with changes in marginal damages. In this paper, we combine probability distributions over MACs with uncertain marginal damages to analyze climate technology policy.

We will address this dynamic decision problem by incorporating the uncertain returns to R&D in a stochastic program. In order to do this, we must first answer two questions: (1) How will different technologies impact the MAC; and (2) What is the probability distribution over different outcomes of technical change? In this section we discuss our prior work, which provides some initial answers to these questions.

3.2 Expert Elicitations on Energy Technology

Past data on technological advance contains little information about future technological breakthroughs. In fact, a technological breakthrough, by its nature, is unique; and therefore we cannot use past data and relative frequencies to construct a probability distribution over success for future breakthroughs. Yet, current decisions depend on understanding the likelihood of such breakthroughs. For example, sound government technology R&D policy should consider the *likelihood of success* and the *impacts of success*, along with the total cost of a program, when making decisions [**?**]. When past data is unavailable or of little use, the alternative is to rely on subjective probability judgements [**?**]. Expert Elicitations are a formal method for gathering these expert judgements.

Decision analytic methods including expert elicitations [**?**] have been applied productively to R&D in numerous industries (automotive, pharmaceutical, electronics, etc. See for example [**?**][**?**][**?**]) as well as issues relating to societal decisions [**?**][**?**][**?**]. A draft of a recent USEPA white paper available on the web says that Expert Elicitation "... should be considered to characterize uncertainty, where it can not be addressed adequately by existing data or additional studies within the necessary timeframe." Most relevantly, a National Research Council study recommends that the U.S. Department of Energy use panel-based probabilistic assessment of R&D programs in making funding decisions [**?**].

In [**?**][**?**][**?**] we have performed expert elicitations on CCS, nuclear, and solar PV technologies. (Also see [**?**] for batteries for vehicles). The products of the elicitations included explicit definitions of endpoints for each technology, and probabilities of achieving those endpoints for given funding trajectories. In Tables 1 - 3 we report the relevant results. The first column identifies the technology category and the second column lists the sub-categories we considered for each technology. The third column gives the Net Present Value (NPV) of the funding trajectory considered. The funding trajectories themselves varied by yearly amount and by the number of years. We have used a discount rate of 5% to calculate the NPVs. We considered multiple funding trajectories for some technologies. The fourth column reports the average probability of success elicited from the experts. In some cases, we had defined two different levels of success. In these cases, the probability on the top is the probability for high success, on the bottom for lower success. For example, Organic solar cells have two levels of success for each funding trajectory; Inorganic solar cells have only one level of success. The fifth and sixth columns represents the impact on the MAC and will be discussed in Section 3.3 below.

Technology	Project	NPV of Funding (000,000)	Probabiltiy of success	Alpha	Shift
Solar	Organic	$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	$39	26.7%	0.022	0.007
		$77	44.3%		
	3rd Generation	$386	2.0%	0.050	0.017

Table 1: Summary of Assessment Results for Solar.

Technology	Project	NPV of Funding (000,000)	Probabiltiy of success	Alpha	Shift
CCS	Pre-Com	$39	2.7%	0.347	0.004
		$154	11.0%		
		$386	22.3%		
	Chem-Loop	$19	8.0%	0.380	0.020
		$38	29.5%		
		$56	42.0%		
	Post-Com	$52	59.0%	0.319	-0.008
		$224	70.0%		
		$519	78.5%		

Table 2: Summary of Assessment Results for CCS.

3.3 Computational MACs using MiniCAM

We then determined how the technologies would impact the MAC, if they achieve the defined endpoints. Specifically, we derived MACs for the year 2050 under different assumptions about technological pathways. We considered all combinations of technological success. The analysis was conducted using the MiniCAM integrated assessment model. MiniCAM is a global model that looks out to 2095 in 15-year timesteps. It is a partial-equilibrium model, with 14 world regions that includes detailed models of land-use and the energy sector. See Brenkert et al. [?] and Edmonds et al. [?] for more discussion of the model. Assumptions for technologies other than the specific technologies we were considering were based on the version of MiniCAM used in the Climate Change Technology Program (CCTP) MiniCAM reference scenario [?]. See [?][?][?] for more detailed discussions of our methods and assumptions on related technologies.

Here we will briefly address some of the difficulties we faced in modeling each of the technologies. First, since solar is an intermittent resource – it cannot be turned off and on – it potentially poses problems for integration onto the electricity grid. The baseline assumption in MiniCAM is that when the penetration of solar into the electricity grid reaches 20%, every additional kW of solar installed requires the installation of a kW of gas-fired backup generation. This is the scenario we have used in the analysis in this paper. We have, however, also modeled the other extreme, simply assuming that there is no problem with grid integration. This would arise if, say, we had free electricity storage. These two scenarios give an envelope of the impact that solar might have, and illustrate the benefits to developing technologies to address grid integration problems.

We faced a range of difficulties in modeling nuclear power. Many of the advantages of new technologies, such as high-temperature reactors and Fast reactors, are not easily modeled or valued. These include a reduction in proliferation concerns, a reduction in radioactive waste, and a reduction in the complexity of the technology. Moreover, the nuclear science experts we worked with gave us relatively low costs for our technological endpoints, of $1500 or $1000/kW; nuclear economists have commented that these costs may be unrealistic. The baseline assumption in MiniCAM is that Light Water Reactors (LWR) will have a cost of $2100/KW in 2020. Our results focus mainly on improved LWR, and should be interpreted as reducing their cost by more than 50% below what otherwise would occur. We do not explicitly model limits to the penetration of nuclear due to political-economy reasons.

Finally, there is concern about the widespread implementation of CCS. The U.S. Department of Energy (DOE) Carbon Sequestration and Technology Roadmap [?] lists a number of challenges, including permanence; monitoring, mitigation and verification; permitting and liability; and public acceptance. Our elicitations did not consider these issues explicitly. An NAS elicitation on this subject [?], however, did consider public opposition based on the risk of sequestration; regulatory issues; and physical siting requirements. They report that the "average panel probability that the large-scale sequestration would be allowed is .66 without DOE's research support and increases to .77 with DOE's support." We use a baseline value of 70% as the likelihood that CCS will be allowed.

Figure 5 presents four representative MACs plus the baseline. Besides the baseline we show the MACs generated assuming (1) success in purely organic solar cells only; (2) success in chemical looping CCS only; (3) success in LWR only; and (4) success in all three of these technologies simultaneously. The left panel show the impacts on low abatement levels and the

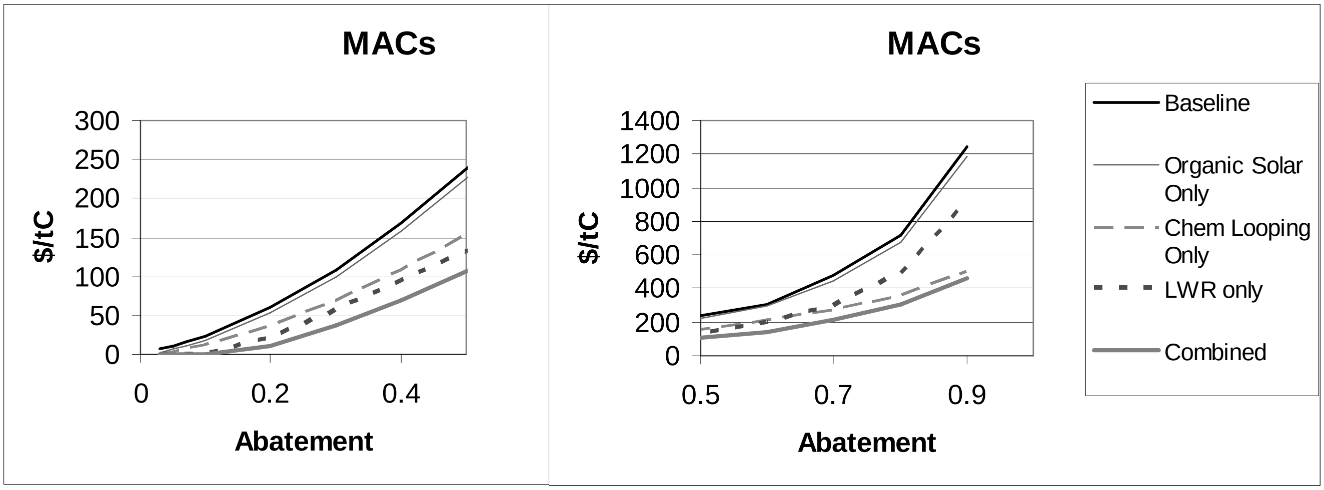

Figure 5: Sample MACs

right panel for high abatement levels. Note that solar only has a small impact on the MAC, even at a cost of 3c/kWh, due to the assumptions about grid integration. Nuclear and CCS have different types of impacts on the MAC. At low abatement levels, nuclear has the greatest impact. In particular, success in LWR implies that carbon emissions would drop by about 10% *even in the absence of a carbon policy.* At high abatement levels, however, CCS begins to dominate, significantly reducing the MAC at abatement levels above 70%. Finally, the combined MAC shows that the technologies are substitutes to a large degree.

If we combine these empirical curves with the elicited probabilities above we have random MACs – a probability distribution over a discrete number of curves. However, working with random functions is challenging theoretically and computationally. So, in the next subsection we parameterize the functions to make them more tractable to work with.

3.4 Parameterization of the MAC

In this section we discuss how we parameterized the impact on the MAC to produce a probability distribution over MACs (in terms of our parameter α) for different levels of funding of different projects. We used the data generated by MiniCAM to estimate a smooth relationship between technical change and the impacts on the MAC. We observed that the effect on the MAC is a combination of a downward pivot and a downward shift. See Figure 6 for an example. Thus, we parameterize the impact by two parameters, α measuring the pivot and h measuring the shift:

$$\widetilde{MAC}\left(\mu;\alpha,h\right)=\left(1-\alpha\right)\left[MAC\left(\mu\right)-h*MAC(0.5)\right] \tag{5}$$

where the tilda represents the MAC after technical change parameterized by α and h and $MAC\left(\cdot\right)$ is the original MAC, before technical change. The first term on the right hand side pivots the MAC down. The second term in the square brackets shifts the MAC downward by a fixed amount. The constant h differs for each individual technology and technology combination. In order to make the parameterization portable to multiple models, we anchored the shift to the

Technology	Project	NPV of Funding (000,000)	Probabiltiy of success	Alpha	Shift
Nuc	LWR	$173	21.3%	0.325	0.118
		$260	33.8%		
		$346	60.0%		
	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
	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 Assessment Results for Nuclear.

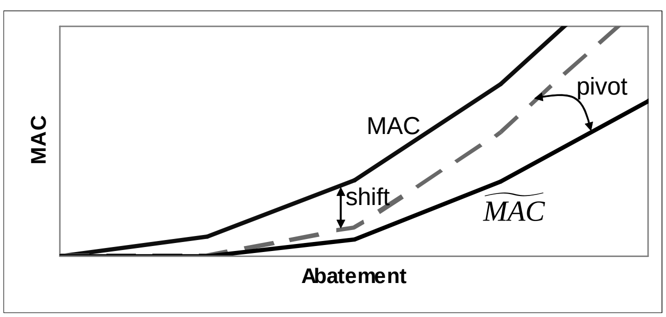

Figure 6: A stylized example of a shift and a pivot to the MAC.

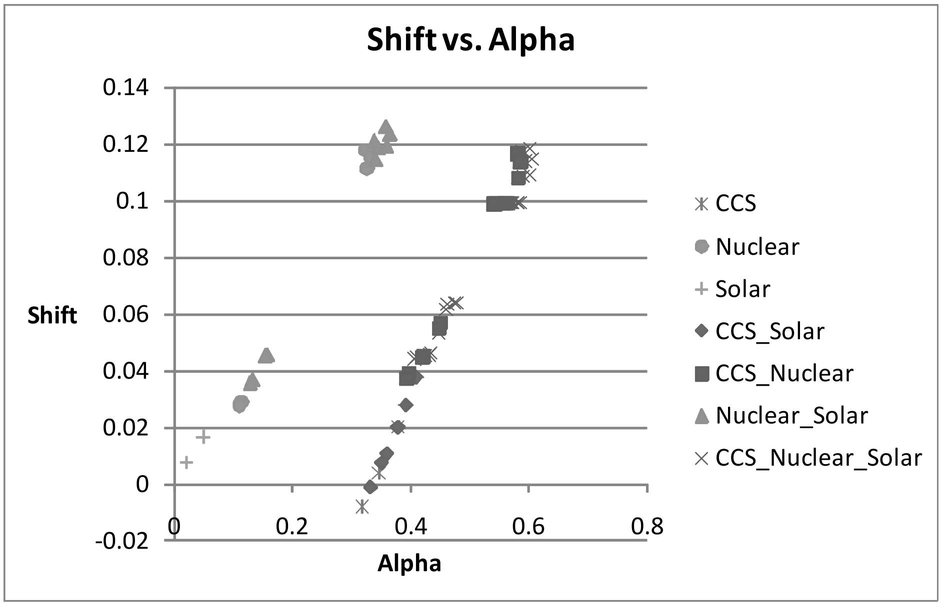

Figure 7: The shift and pivot of all technology combinations

marginal cost of 50% abatement. For the individual technologies, we simply estimated the values for α and h from the empirical MAC curves using a least squares method. For the combined technologies, we assumed that the pivot would be a multiplicative combination of the single technologies: $\alpha_{CSN} = 1 - (1 - \alpha_C)(1 - \alpha_S)(1 - \alpha_N)$. We then estimated the values for h using the least square method. The values of α and h for each single technology are given in Tables 1 - 3. Figure 7 graphs the values of h and α for each individual technology and technology combinations. Nuclear, Solar, and their combinations have relatively weaker pivots and stronger shifts than portfolios that include CCS. This matches what can be seen in Figure 5. CCS has mostly a pivot effect, with virtually no impact when the carbon price is very low, and a strong impact when it is high. Nuclear, on the other hand, shifts the MAC downward, but has a lower pivot effect, as seen from the right panel.

4 The Portfolio Model

Figure 8 represents our model in the form of an Influence Diagram. The square nodes represent decisions. The first stage decision, our focus, is how to allocate R&D investment across a number of projects. The oval nodes represent uncertainties. We are uncertain about which technologies will be successful or not. The arrow between the decision node and the uncertainty node means that the choice of which programs to fund influences the probability distribution over successful technologies. The double-lined oval node is a deterministic node. It means that we assume that the cost of abatement can be derived if we know which set of technologies is available. The oval in the lower part of the figure indicates that climate damages are also uncertain. The arrows pointing to the second stage decision, how much to abate, indicate that this decision is made after learning has taken place. The future decision about how much to abate will be made based

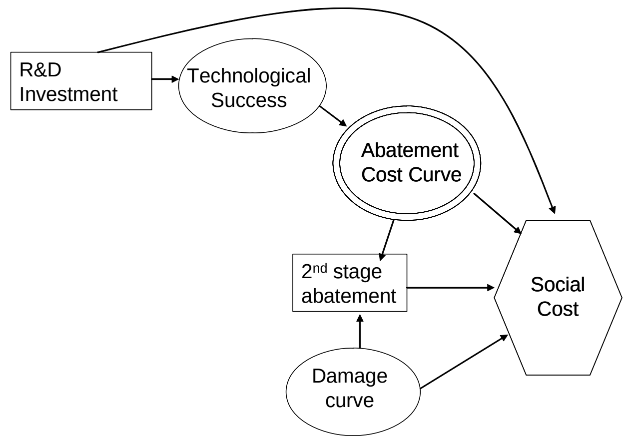

Figure 8: Influence Diagram of R&D Problem

on knowledge about the set of technologies available and about climate damages.

The overall goal of the model is to minimize the sum of expected abatement costs and expected damages for a given R&D budget. The first stage decision is *which set of R&D projects to fund*. For this paper we will focus on three categories – PV, nuclear fusion, and CCS. Within each category we consider 3-4 projects based on specific sub-technologies (such as Purely Organic Solar Cells, Light Water Reactors, or Post-combustion separation), and 2-3 funding levels for each sub-technology, as laid out in Tables 1 - 3. Each potential funded portfolio leads to a probability distribution over *successful* portfolios (based on our expert elicitations). Each successful portfolio will determine a Marginal Abatement Cost Curve, based on our parameterizations as described above.

Climate change damages are also uncertain. We calibrate a baseline damage function, as a function of abatement, to the DICE-2007 model [?]. We develop two- and three-point probability distributions over the damages using estimates from Nordhaus [?]. Part of our analysis is to perform sensitivity analysis over these probability distributions to understand the role of increasing risk.

The second-stage decision is *how much to abate*, for a given damage and abatement cost curve. In the absence of a corner point, abatement will be chosen so that the marginal cost of abatement, after technical change, is equal to the marginal damages. We will consider corner points where the marginal cost of abatement is less than marginal damages and full abatement is optimal.

There are two main approaches to modeling and solving a two-stage problem like this, namely stochastic programming and dynamic programming. Our problem presents challenges for both of these approaches. For traditional dynamic programming or decision trees the imposition of constraints and a large number of choices leads to a problem of intractable size. In our case, even though we consider only three categories of technology, a decision tree based on this problem would have about 250 million endpoints, with each endpoint a non-linear optimization problem. Stochastic programming, on the other hand, allows us to apply convex optimization methods to solve the problem numerically. However, classical stochastic programming is restricted to problems in which the probability distributions of the random parameters are fixed and do not depend on the decisions taken. The natural structure of our problem (and indeed most R&D problems) is that investment in a particular technology increases the probability of success in that technology. (See [**?**] in which probability of success is a function of investment). In particular, our experts have given us probabilities conditional on funding trajectories.

We approach this problem using stochastic programming, and develop methods to deal with the challenges described above. To this end, we first describe the basic formulation of the problem, in which probability of success is a function of investment. This assumption, as well as other parts of the problem, lead to a highly non-convex objective function. In the appendix we describe a procedure to reformulate and solve the problem as an equivalent convex problem.

4.1 Initial Non-convex model

We let the indices i and j represent the technology category (solar, CCS, nuclear) and the specific project within the category, respectively. Further, the index $k = l, h$ represents the investment level. The key integer decision variables are x_{ijk}, which equal 0 if there is no investment in project ij at funding level k, and 1 otherwise. The second stage continuous decision variable is abatement $\mu \in [0...1]$, i.e. the fraction of emissions reduced below a business-as-usual level. This variable is conditioned on the state of climate damages, represented by a random multiplier Z; and by the state of the invested technologies, represented by the random vector $\overrightarrow{\alpha}$. The objective is to minimize the sum of abatement costs and damage costs as follows:

$$\min_{x,\mu\left(\overrightarrow{\alpha},Z\right)} E\left[c\left(\mu;\overrightarrow{\alpha}\right) + ZD\left(\mu\right)\right] \tag{6}$$

Note that the investment in a technology is made without information on technical success or climate damages, while abatement is chosen conditional on technical success and damages, i.e. it is a second stage decision. We are constrained by the R&D budget B, and by the fact that we can only invest in a project at one level:

$$\sum_i \sum_j \sum_k f_{ijk} x_{ijk} \leq B \tag{7a}$$

$$\sum_k x_{ijk} \leq 1, \quad \forall i, j \tag{7b}$$

where f_{ijk} is the required level of investment for funding level k of project ij.

We assume that the probability of technical success in any technology is independent of other technologies (and of the damages of climate change). Thus, the probability of any realization of the random vector $\overrightarrow{\alpha}$ is simply the product of the probability of the individual parts of that realization.

According to our elicitation, however, the probabilities of success for individual projects depend on whether that project has been invested in or not, as seen in Tables 1-3 above. In parallel with the general stochastic modeling framework, we perform the following steps to define the input distributions of the model. *First*, we calculate the probability of each realization of $\overrightarrow{\alpha}$ exogenously, using the probability of success if funded. For each funded project ijk there are three potential outcomes, failure, moderate success, or high success. We index these by $l = -1, 0, 1$. Then, for example, the probability of the event that there is high funding, i.e.

$k = 2$, in organic solar cells ($i = S, j = 1$) and that we get high success in organic cells and no success in anything else is:

$$p_{S12,1} * \prod_{(i,j,k)\neq(S,1,2)} p_{ijk,-1} \quad (8)$$

where $p_{ijk,l}$ represents the probability that funded project ijk will result in outcome l.

Second, we define the outcomes so that they correctly match with the probabilities. The outcome of each realization of $\overrightarrow{\alpha}$ is a vector with entries α_i, $i = S, C, N$, representing the amount of technical change in each category. We assume that only the best technology project in each category will diffuse in the economy. For example, if all solar projects are successful, we assume that the lowest cost technology will take over the market, giving solar a cost of 2.9 cents/kWh and $\alpha_S = 0.045$. Let $\overrightarrow{\omega}$ be the state of the world, a vector containing the realizations l of each project. Then we define the components of $\overrightarrow{\alpha}$ vector as follows:

$$\alpha_i(x; \omega) = \max_{j,k}\{x_{ijk}\alpha_{ijk,l}\} \quad (9)$$

where $\alpha_{ijk,l}$ is a parameter taken from Tables 1-3 above. In this formulation, if we do not invest in technology ij, then $x_{ijk} = 0$. For example, consider again the event that there is high funding in organic solar cells and that we get high success in organic cells and no success in anything else. The realization of $\overrightarrow{\alpha}$ associated with this event depends on whether organic solar is funded at the high funding level or not. If $x_{S12} = 0$ then the outcome will be $\overrightarrow{\alpha} = (0, 0, 0)$; if $x_{S12} = 1$ the outcome will be $\overrightarrow{\alpha} = (.045, 0, 0)$. The outcome depends on the decision variable x_{ijk}, while the probability does not.

Between the technology categories, we assume that the pivots are multiplicative, but that the shifts are defined according to dependency relationships between the technologies, as mentioned in Section 3.4 above. Based on the relationships established through simulations, the total shift in the MAC, h can then be defined as:

$$h = K(\mathbf{x}, \overrightarrow{\alpha}) \quad (10)$$

where $K(\mathbf{x}, \overrightarrow{\alpha})$ represents the constant shift value reported in Tables 1-3 for the individual technologies and represented in Figure 7 for combinations. For each possible combination of α values, such mappings can be generated exogenously and included in an optimization model. This is further discussed in Section 4.2.

Given h, the cost is then:

$$c\left(\mu; \overrightarrow{\alpha}\right) = \prod_i \left(1 - \alpha_i\right)\left[c\left(\mu\right) - hc\left(0.5\right)\mu\right] \quad (11)$$

where $c\left(\mu\right)$ is the cost before technical change. Notice that the shift is multiplied by μ; this is because the parameterization above was done on the MAC and now we are working with the cost. We have based our baseline cost on the DICE 2007 model [**?**]:

$$c\left(\mu\right) = b_0\mu^{b_1} \quad (12)$$

The damage function is assumed to be quadratic:

$$M_0(S - M_1\mu)^2 \quad (13)$$

$Zh:$	1(no risk)	3(medium risk)	14.6(high risk)	14.6 (baseline)	3 (high damage no risk)	14.6 (high damage high risk)
$P[Z=0]$	-	2/3	.931	.245	–	.795
$P[Z=1]$	-	-	-	.737	–	-
$P[Z=Zh]$	1	1/3	.068	.018	1	.205
Optimal abatement if $Z=Zh$	46%	80%	100%	-	-	-

Table 4: Damage Uncertainty

We calibrated b_0, b_1, M_0, M_1 and S to DICE 2007. The stock of emissions in the atmosphere S is set equal to stock of emissions in 2185 under the Business As Usual (BAU) scenario in DICE, equal to 2.5 trillion metric tons of carbon. The damage constants M_0, M_1 are set so that our damages equal the Net Present Value of damages between 2005 and 2185 in DICE under the BAU and "optimal" scenarios. We used the BAU scenario to calculate that $M_0 = 2.74$. We take the optimal level of abatement (with no technical change) to be the average of the optimal abatement in DICE 2007 over the period 2005 to 2185, or 0.46. Given this, $M_1 = 0.597$. The value of $b_1 = 2.8$, the value in DICE. We set b_0 so that the optimal abatement is 0.46. This leads to a value of $b_0 = 10.4$.

We consider multiple cases for uncertainty over climate damages, represented in Table 4. High damages, where $Z = 14.6$, are equivalent to a 20% loss in GDP given a 2.5°C increase in mean temperature. Each risk scenario in columns 2 - 5 has a mean of 1. The High Risk case has the highest possible probability for the high damages without allowing negative damages (i.e. benefits). The medium risk case is a Mean-Preserving Spreads (MPS) of the no risk case (See [?] for a definition and discussion of MPS). The High risk case is an MPS of both the no-risk and medium risk case. The last two columns have a higher mean of $Z = 3$.

4.2 Stochastic Programming Formulation

For the formulation of the problem as a two-stage stochastic programming problem, we first expand our definition of ω and let $\omega \in \Omega$ represent a scenario consisting of possible values of the parameters $\alpha_{ijk,l}$ and Z, and define p^ω as the probability of the scenario ω, calculated as described in Section 4.1. Since the scenario definition involves both the vector $\overrightarrow{\alpha}$ and the random parameter Z, we refer to the realized value $\alpha_{ijk,l}$ as α^ω_{ijk} for consistency in the description of the formulation. Note, our convention is that realizations of random variables have ω as a superscript; whereas decision variables that are conditional on the realization have ω as a subscript. The overall stochastic optimization problem can then be expressed as follows,

$$\min_{\mathbf{x}\in\mathcal{X}} \sum_{\omega\in\Omega} p^\omega \{\prod_i (1 - \max_{j,k}\{\alpha^\omega_{ijk} x_{ijk}\})(b_0 \mu_\omega^{b_1} - c_{0.5} h_\omega \mu_\omega) + Z^\omega M_0 (S - M_1 \mu_\omega)^2\} \tag{14}$$

$$\text{s.t. } h_\omega = K(\mathbf{x}, \overrightarrow{\alpha}) \quad \forall \omega \tag{15}$$

$$0 \le \mu_\omega, h_\omega \le 1 \quad \forall \omega \tag{16}$$

where $\mathcal{X}$ represents the set of feasible investment decisions, as defined by (7). Note that problem (14)-(16) is the deterministic equivalent of the stochastic optimization problem (6). On the other hand, the multiplicative nature of the pivot terms in the cost function, i.e. the product $\prod_i (1 - \max_{j,k}\{\alpha^{\omega}_{ijk} x_{ijk}\})$, results in the model being highly nonconvex. Thus, convex optimization based approaches are not applicable to the model, and a convex approximation or reformulation approach is necessary. The nonconvex product term is an integral part of the overall model, and results in a set of bilinear and trilinear components, approximation of which are typically not tight. However, we show in the appendix that an equivalent convex reformulation of the problem can be developed by defining some new variables and revising the definition of some parameters. See the appendix for a detailed description of the model and the solution procedure.

5 Results

Our analysis of the optimal climate change energy technology portfolio under different configurations resulted in several interesting implications from a policy perspective. As part of our analysis, we first considered different R&D budget levels and observed the impact of damage uncertainty on the composition of the optimal portfolio. Then, we investigated the impact of risk and of assumptions about opportunity costs on the overall optimal investment in energy technology R&D in the presence of climate change. We summarize our findings in the next two subsections.

5.1 Composition of Optimal Energy Technology R&D Portfolio

We considered different R&D budget levels and observed the impact of damage uncertainty on the value and composition of the optimal portfolio. The first result is that the optimal portfolio was very robust to different levels of damage risk. Figure 9 shows the composition of the optimal portfolio at budget levels ranging between \$200 and \$2000 million. These portfolios did not change under any of the scenarios in Table 4. We know from previous research that damage risk can impact the optimal investment in technology [?][?]. Thus, this result shows the value of incorporating actual data. Specifically, in this case, the data has lead to projects that are fairly differentiated – some projects (such as chemical looping and LWR) have high probabilities and high payoffs, and therefore get funded regardless of risk, and even regardless of the mean of damages. Additionally, the problem is somewhat sparse, and so there are not a large number of possible portfolios at different budget levels. In future work we will perform extensive sensitivity analysis to see if this results holds under different parameters. If it does hold, it is good news that the optimal R&D investment is robust to uncertainty in climate damages.

Second, we see the effects of this being a "knapsack" problem. We see that solar, in particular, goes in and out of the portfolio at different budget levels. The solar projects (under our assumption of grid integration limits) are less efficient than some of the other projects, but also less costly. Thus, for example, we see a significant investment in solar at the \$200 million budget level; but this investment is reduced in favor of nuclear when the budget increases. We do see strong diversification – all three technology categories come in to the optimal portfolio even at a fairly low budget. At higher budget levels, not shown here, nuclear dominates the portfolio.

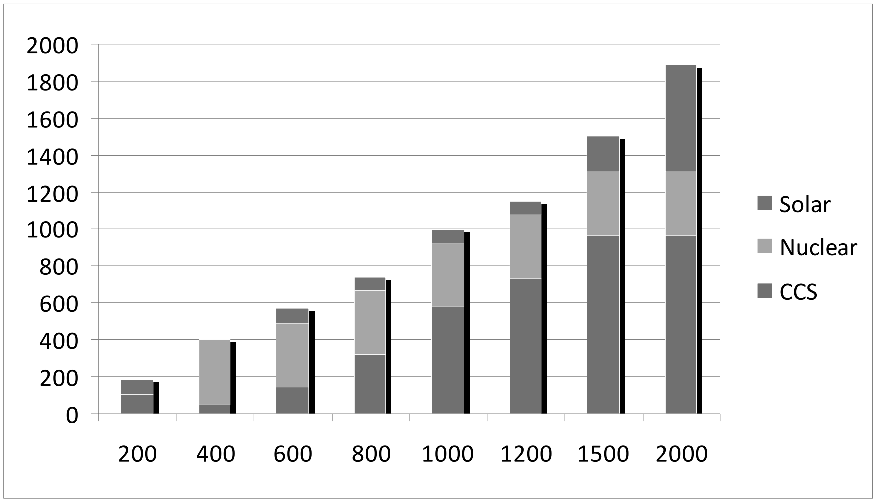

Figure 9: Optimal Portfolios

Figure 10 shows how the expected total social cost (damages plus abatement) is impacted by R&D investment, in the four risk cases in columns 2-5 in Table 4. The curves in the figure are normalized so that all cases appear on the same scale as the no-risk case.[10] In addition, Table 5 shows the non-normalized approximations of the marginal value of R&D for each budget level. The table shows the additional value of the portfolio divided by the amount of R&D investment, in billions. Notice that R&D is very efficient. The lowest budget we considered has an NPV of $0.2 billion and reduces the expected total social cost by over $300 billion in the no-risk case, a marginal value of $1,632 for every dollar spent. An additional $0.4 billion investment leads to an additional $370 billion reduction in costs. Even in the high risk case, the additional $0.4 billion reduces costs by about $120 billion, leading to a marginal value of $403 for every dollar spent. Note that there is an "elbow point" in each of the graph lines, where the cost savings from a bigger portfolio slows down considerably. This happens at a budget of $600 million, and consists of a portfolio including a high investment in chemical looping, LWR, and purely inorganic PVs, along with medium level investments in the other two CCS technologies.

We pointed out above that the composition of the optimal portfolio at given budget levels is constant over a variety of different risk configurations. However, Figure 10 and Table 5 show that the *value* of R&D is impacted by the level of risk. First, R&D has the least value in the high risk case. This is because in that scenario we either have no damages and no abatement, or we have very high damages that lead to full abatement regardless of the technology. Thus, the technology reduces the cost of abatement, but does not change the optimal level of abatement – it has no environmental-side effect. As a contrast, in the no risk case, the presence of technology not only lowers the costs of abatement for a given level of abatement, but also leads to optimally lower abatement. In fact, when $Z = 1$ our results show that the overall expected cost of abatement increases as the R&D budget increases – the optimal level of abatement increases enough that it outweighs the reduced cost of abating any given level. That is, the technology has a significant

[10] Total Expected Cost is lower under risk, since abatement is increased when damages are high. See [**?**] for a discussion of this.

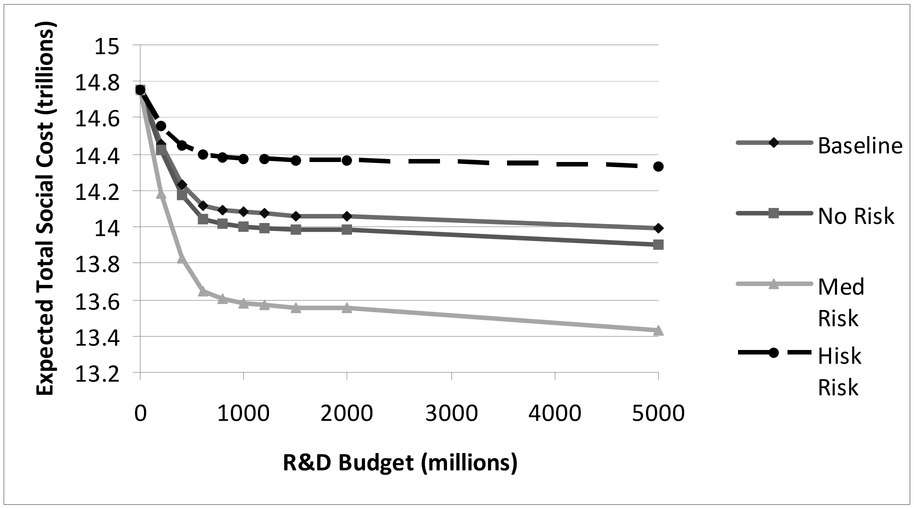

Figure 10: Expected Total Social Cost

	Marginal Value ($ billions)			
Budget ($ millions)	No Risk	Baseline	Medium Risk	High Risk
200	1632	1388	2508	707
400	1247	1024	1572	403
600	660	536	822	191
800	119	98	162	38
1000	81	67	112	27
1200	42	34	53	12
1500	37	30	47	11
2000	3	2	3	1
5000	28	21	35	8

Table 5: Marginal value of R&D at different budget levels

environmental-side benefit. Thus, it has overall more value.

We see, however, that the value of R&D is non-monotonic in risk, increasing significantly in the medium risk case. This is because, when $Z = 3$, we get both cost-side and environmental-side benefits. In this case, our results show that both expected damages and the overall cost of abatement decrease at higher budget levels.

Figure 11 illustrates this point. The figures show the expected MAC when the budget is $600 million as well as the baseline MAC; and the marginal damages when $Z = 1, 3$. The left-hand chart shows the impact of technical change when $Z = 1$. Optimal abatement increases from 46% to about 65%, thus there is environmental-side benefit. The total cost of abatement is the area under the curve. It can be seen that the total abatement cost in this case is slightly higher after technical change. The right-hand panel shows the impact of technical change when Z=3. Optimal abatement increases from 80% to 100%, thus again there is environmental side benefit. Overall abatement cost also decreases in this case, as can be seen by comparing the

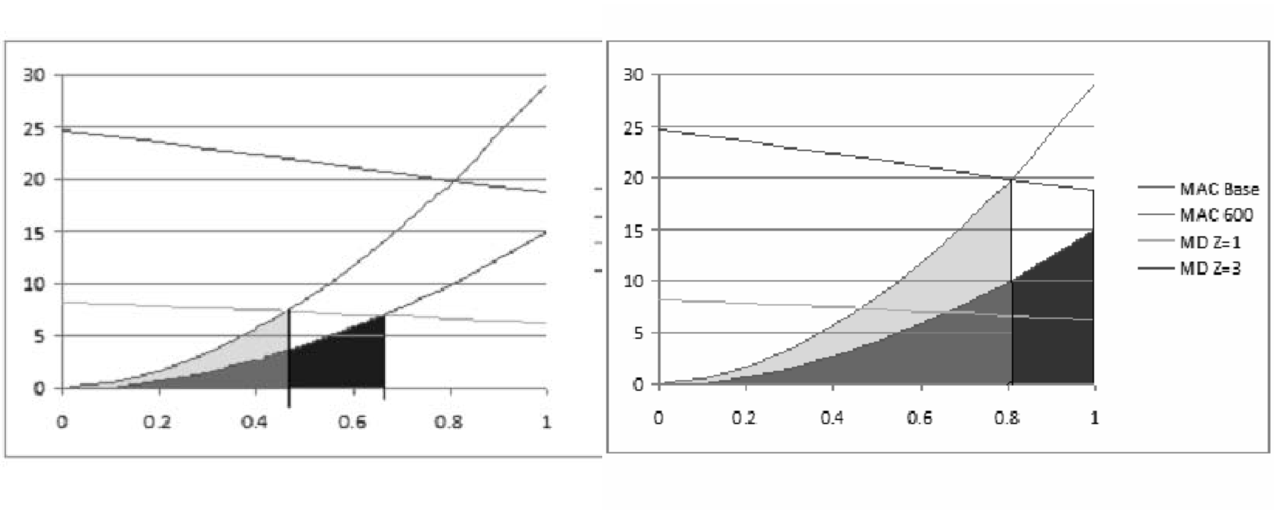

Figure 11: Optimal Abatement and Total Cost of Abatement

lightest wedge (cost saved after technical change) with the darkest trapezoid (costs added after technical change because of higher abatement). Thus, overall, technical change has more value in the second case than the first.

Note that we will observe a similar phenomenon if we have a fixed target for abatement (or equivalently a fixed concentration target). In these cases, technology will lower the costs of hitting the fixed target, but will not have an impact on the target. Thus, the focus on a concentration target (which is very common in the climate change literature), rather than an optimal level of abatement, reduces the value of R&D to society. If we fail to recognize that we can adjust emissions based on the outcome of technical change, we will systematically under-value investments in R&D.

5.2 Optimal Level of R&D Investment

In this section we calculate the overall optimal investment in R&D under different assumptions about the actual cost of R&D. In Tables 1-3 we report the NPV of the funding levels from earlier expert elicitations [?][?][?]. These funding trajectories represented the amount of money going into the hands of high quality researchers in the appropriate areas. The funding trajectories do not account for administrative costs of awarding the funding, nor do they account for the possibility of "pork" – money awarded by earmark for political reasons rather than based on scientific merit. Moreover, money spent on R&D is considered to have a particularly high opportunity cost in the economy, perhaps up to 4 times as much as the out-of-pocket expense [?][?].

Although exact nature and amount of these opportunity costs are still an open question, we perform an analysis over a range of opportunity costs. We show results for the cases of no opportunity costs as well as total costs equal to 2, 4, and 8 times the net costs. The lower assumption would hold if "pork" was minimal and only about 50% of new energy R&D was replacing other kinds of R&D (see [?]), while the highest assumption would hold if "pork" doubled the cost of R&D and all energy R&D replaced other kinds of R&D. We find the optimal portfolio under three risk cases and these four assumptions about cost. The results are shown in Tables 6 - 8. Columns 2-10 in each table show the optimal investment in each specific technology; column 11 shows the overall optimal net investment in R&D (that is, not including opportunity costs); and the last column shows the total expected social cost (including the opportunity cost of investment).

Coef	Investments ($ million)									Tot. Inv.	Tot. Cost
	Pre C	Chem L	Post C	LWR	HTR	FR	Org.	Inorg.	3rd g	($bil)	($ tri)
1	386	56	519	346	3089	15443	830	77	386	21.132	13.86313
2	386	56	519	346	3089	15443	830	77	386	21.132	13.88426
4	386	56	519	346	3089	0	830	77	0	5.303	13.91521
8	386	56	519	346	3089	0	116	77	0	4.589	13.93571

Table 6: Optimal portfolio as a function of the opportunity cost multiplier for no risk

Coef	Investments ($ million)									Tot. Inv.	Tot. Cost
	Pre C	Chem L	Post C	LWR	HTR	FR	Org.	Inorg.	3rd g	($bil)	($ tri)
1	386	56	519	346	3089	15443	830	77	386	21.132	11.84713
2	386	56	519	346	3089	15443	830	77	386	21.132	11.86826
4	386	56	519	346	3089	4633	830	77	0	9.936	11.90974
8	386	56	519	346	3089	0	830	77	0	5.303	11.93342

Table 7: Optimal portfolio as a function of the opportunity cost multiplier for medium risk

Coef	Investments ($ million)									Tot. Inv.	Tot. Cost
	Pre C	Chem L	Post C	LWR	HTR	FR	Org.	Inorg.	3rd g	($bil)	($ tri)
1	386	56	519	346	3089	4633	830	77	0	9.936	10.33594
2	386	56	519	346	3089	0	116	77	0	4.589	10.34218
4	386	56	519	346	3089	0	116	77	0	4.589	10.35316
8	386	56	519	346	1544	0	116	77	0	3.044	10.36735

Table 8: Optimal portfolio as a function of the opportunity cost multiplier for high risk

It can be seen from the three tables that the optimal investment level varies in the risk of climate damages. These results are summarized in Figure **??** . The pattern that emerges is consistent with the results in Figure 11, in which R&D has the highest value in the medium risk case and the lowest value in the high risk case. We see here that the optimal net investment in R&D is highest in the medium risk case and lowest in the high risk case. Notice that when the opportunity cost is low, the entire portfolio is funded under the no- and medium-risk cases.

Consistent with our findings that the composition of the portfolio is robust to risk, it appears that the value of the individual technologies is not strongly effected by risk. If we read each table from top to bottom, we can see which technologies get reduced funding or leave the optimal portfolio as the opportunity cost gets higher. It appears that the first technology to be reduced is 3rd generation solar, followed by the Feeder Reactor, followed by organic solar cells, and finally the HTR reactor.

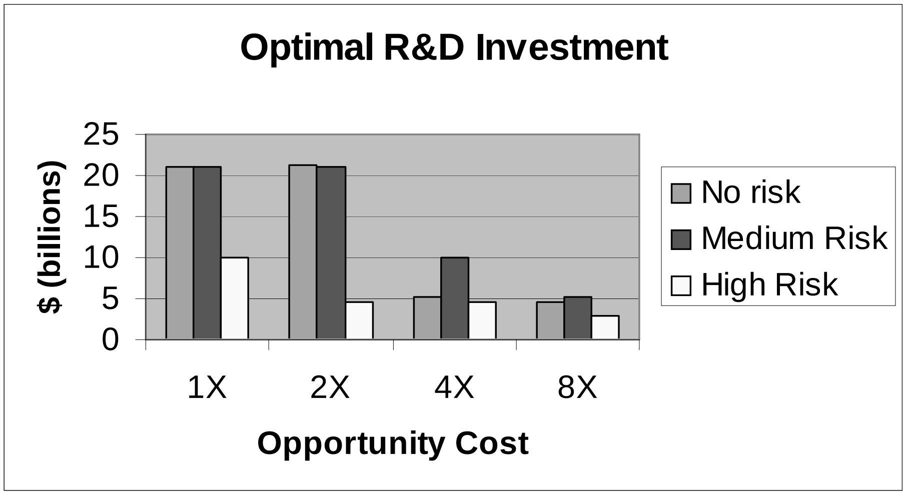

Figure 12: Optimal R&D Investment

6 Conclusions

In this paper we have gone beyond theoretical analysis to present results from a data-based climate change energy technology R&D portfolio model. Our R&D portfolio model has provided a number of insights. While it is easy to show theoretically that the optimal portfolio can depend on the level of risk, we have found in our data-based model that the optimal portfolio (conditioned on specific budget levels) is robust to climate damage risk. This is good news, since determining the probability distribution over climate change damages is very difficult. Moreover, the optimal portfolio is even robust to the mean of the distribution. Second, we do see a high level of diversification, with even less-promising technologies included in the portfolio, although this is partly a result of it being a knapsack problem.

Third, while the portfolio at any given budget level is robust to risk, this is not true for the value of R&D. This leads to our next result, which is that the optimal level of spending depends explicitly on the probability distribution around climate damages. Fourth, we see that R&D and technical change has less value when future emissions levels are fixed, such as is the case for very high risk, or for fixed emissions targets. CCS in particular seems to have the most value when emissions are flexible. Related to this point, very high risk – when there is a chance of truly catastrophic damages that will induce full abatement – will favor technologies that reduce total cost (versus technologies that reduce the MAC). Finally, the value of technology is non-monotonic in risk, with the maximum value being in cases where technology leads to higher abatement and significant reductions in abatement costs.

There is a great deal of work left to be done. The results depend on the subjective probabilities collected in expert elicitations. As we mention in Section 3.3 above, there are a number of concerns about the data and the assumptions we have had to made in order to implement the results in an economic model. Therefore, we plan to perform extensive sensitivity analysis to determine which data points are most crucial to the results. Together with this, we can perform an analysis on the value of information, in order to inform future research projects of this kind.

We would also like to include a larger number of technologies. We are in the middle of doing the economic analysis needed to include batteries for vehicles and biofuels into the analysis. Finally, the results indicate that the optimal level of R&D spending may depend on the damage risk scenario. We plan on testing this possibility, under different assumptions about the opportunity costs of R&D spending.

Acknowledgement 2 *This research was partially supported by the Office of Science (BER) U.S. Department of Energy, Grant No.DE-FG02-06ER64203 and by NSF under award number SES-0745161. We thank Yiming Peng for research assistance.*

A Appendix

A.1 Equivalent Convex Reformulation

To develop an equivalent convex formulation, we first let ϕ_i be a nonnegative variable such that it is equal to the value of $-\ln(1 - \max_{j,k}\{\alpha_{ijk}\})x_{ijk}$ for $j, k \in \arg\max_{j,k}\{\alpha_{ijk}\}$. Note that these variables are defined for each scenario, but we leave out the index ω in these definitions for the clarity of presentation. Further, we define a new nonnegative variable $w = h + \mu$, and binary indicator variables δ_{ijk} and β_i to represent the modified problem structure. β_i corresponds to the case with no investment in technology i, while δ_{ijk} is an auxiliary variable used to indicate whether the corresponding set of constraints holds in the model. Further, for technology category i, δ_{ijk} identifies the funded project determining the value of α_i, which is the highest realized value among all funded project returns in that category. In addition, we let the random parameter $\bar{\alpha}^{\omega}_{ijk}$ represent $\ln(1-\alpha^{\omega}_{ijk})$, which is calculated exogenously. Finally, we define the set of variables $y^{\pi}_{i,i',i''}$ for all $i, i', i'' \in C, N, S$, where π corresponds to a distinct combination of possible α_{ijk} values for the three technology categories. The variables y are used to denote the dependency relationships that apply to the shift parameter h in a given solution to the problem. We will refer to the combined set of y variables as y^{π}_{I}, and assume that a constant K^{π}_{I} is calculated exogenously for each possible combination.

With these definitions and modifications, the following equivalent formulation of the climate change energy technology R&D problem can be developed:

$$\text{Minimize} \quad \sum_{\omega} p^{\omega}[(e^{-\sum_i \phi_{i\omega}+\ln(b_0\mu_{\omega}^{b_1}-\frac{1}{2}c(0.5)(w_{\omega}^2-h_{\omega}^2-\mu_{\omega}^2))}) + Z^{\omega}M_0(S-M_1\mu_{\omega})^2] \tag{17}$$

$$\text{subject to} \quad \sum_i \sum_j \sum_k f_{ijk}x_{ijk} \leq B \tag{18}$$

$$\sum_k x_{ijk} \leq 1, \quad \forall i, j \tag{19}$$

$$\phi_{i\omega} + \bar{\alpha}_{ijk}^{\omega} x_{ijk} + M\delta_{ijk\omega} \leq M \quad \forall i, j, k, \omega \tag{20}$$

$$\phi_{i\omega} + \bar{\alpha}_{ijk}^{\omega} x_{ijk} + m\delta_{ijk\omega} \geq m \quad \forall i, j, k, \omega \tag{21}$$

$$\sum_j \sum_k \delta_{ijk\omega} + \beta_i = 1 \quad \forall i, \omega \tag{22}$$

$$\phi_{i\omega} + M\beta_i \leq 1 \quad \forall i, \omega \tag{23}$$

$$h_{\omega} - \sum_I y_{I\omega}^{\pi} K_I^{\pi} = 0 \quad \forall \omega \tag{24}$$

$$y_{I\omega}^{\pi} = 1 \Leftrightarrow \sum_{i \in I}(\sum_j \sum_k \delta_{ijk\omega}\alpha_{ijk}^{\omega}) = \alpha_I^{\pi} \quad \forall I, \pi \tag{25}$$

$$w_{\omega} = h_{\omega} + \mu_{\omega} \quad \forall \omega \tag{26}$$

$$\delta_{ijk\omega} - x_{ijk} \leq 0 \quad \forall i, j, k, \omega \tag{27}$$

$$\mathbf{x}, \mathbf{y}, \delta, \beta \in \{0, 1\} \tag{28}$$

$$0 \leq \mu, h \leq 1; w, \phi \geq 0. \tag{29}$$

where α_I^{π} refers to the sum of the α values for the combination π, and M and m are upper and lower bounds based on the corresponding constraints. The objective function (17) in the above formulation is based on two reformulation steps. First, the bilinear term $h\mu$ is expressed as a function of the new variable w, as by definition $w^2 = h^2 + 2h\mu + \mu^2$. Then, we describe the product terms using the corresponding natural logs. The constraints (18) and (19) are the first stage constraints (7). The inequalities (20) and (21) ensure that the value of ϕ_i is equal to $\bar{\alpha}_{ijk}x_{ijk}$ if project ijk is selected and $j, k \in \arg\max_{j,k}\{\alpha_{ijk}\}$, while (22) is used to define β_i such that it will be 1 if no investment is made in technology i. Similarly, (23) ensures that $\phi_i = 0$, if no investment is made in the technology. Based on exogenous parameters K_I^{π}, constraints (24)-(25) define the variable h as described in (15). The relationships enforced through constraint set (25) are not stated explicitly for the sake of clarity, but these relations are modeled using standard integer programming methods. Constraints (26) define the variable w, and finally the inequality (27) ensures that a project can contribute to the portfolio only if it is selected.

Problem (17)-(29) is an integer linear program with a nonlinear objective function and linear constraints. Further the objective function is convex as we show below:

Theorem 3 *Problem (17)-(29) is convex.*

Proof: Since the problem contains linear constraints, it suffices to show that the objective function (17) is convex in the decision variables. Note that this function consists of two components, an exponential term and a quadratic function of the variable μ. It is trivial to show that the quadratic component is convex.

For the exponential term, we know that the exponentiation of a convex function is convex. Thus, the problem reduces to showing that $g(h_\omega,\mu_\omega) = -\ln(b_0\mu_\omega^{b_1} - \frac{1}{2}c(0.5)(w_\omega^2 - h_\omega^2 - \mu_\omega^2)) = -\ln(b_0\mu_\omega^{b_1} - \frac{1}{2}c(0.5)((h_\omega+\mu_\omega)^2 - h_\omega^2 - \mu_\omega^2))$ is convex. Note that $g(h_\omega,\mu_\omega)$ is twice differentiable, and the Hessian $H_g(h_\omega,\mu_\omega)$ is given by:

$$\begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}$$

where we use the values listed in Section 4.1 for b_0, b_1 and $c(0.5)$ to obtain

$$a_{11} = \frac{2.22\mu_\omega^2}{(10.4\mu_\omega^{2.8} - 0.745(h_\omega+\mu_\omega)^2 + 0.745h_\omega^2 + 0.745\mu_\omega^2)^2}$$
$$a_{12} = a_{21} = \frac{1.49(10.4\mu_\omega^{2.8} - 0.745(h_\omega+\mu_\omega)^2 + 0.745h_\omega^2 + 0.745\mu_\omega^2) - \mu_\omega(43.39\mu_\omega^{1.8} - 2.22h_\omega)}{(10.4\mu_\omega^{2.8} - 0.745(h_\omega+\mu_\omega)^2 + 0.745h_\omega^2 + 0.745\mu_\omega^2)^2}$$
$$a_{22} = \frac{(29.12\mu_\omega^{1.8} - 1.49h_\omega)^2 - 52.42\mu_\omega^{0.8}(10.4\mu_\omega^{2.8} - 0.745(h_\omega+\mu_\omega)^2 + 0.745h_\omega^2 + 0.745\mu_\omega^2)}{(10.4\mu_\omega^{2.8} - 0.745(h_\omega+\mu_\omega)^2 + 0.745h_\omega^2 + 0.745\mu_\omega^2)^2}$$

Clearly, $a_{11} \geq 0$, as all of its components are nonnegative. Further, it can be shown through algebraic manipulation that $a_{22} \geq 0$ holds for the ranges $0 < \mu_\omega \leq 1$ and $0 < h_\omega \leq 1$. Similarly, $|H_g(\mu_\omega,h_\omega)| \geq 0$, as the determinant of the matrix is given by

$$\frac{61.92\mu_\omega^{4.8}(\frac{0.08h_\omega^2}{\mu_\omega^{2.8}} - \frac{0.31h_\omega}{\mu_\omega} - 1.71\mu_\omega^{0.8})}{(1.49h_\omega\mu_\omega - 10.4\mu_\omega^{2.8})^4} \tag{30}$$

Hence, $H_g(\mu_\omega,h_\omega)$ is positive semidefinite, and $g(h_\omega,\mu_\omega)$ is convex. It follows that problem (17)-(29) is convex.

Given the above result, the problem (17)-(29) can be solved using any nonlinear integer programming solver or through a branch and bound implementation, provided that the number of considered scenarios is not large. For large number of scenarios, which is the case for the climate change energy technology portfolio model, sampling based procedures based on solving randomly sampled small scale instances can be used to determine good or near-optimal solutions, which we describe in the next subsection.

A.2 Solution Approach

We make use of the sample average approximation (SAA) method, a Monte Carlo sampling technique that approximates a stochastic program by a smaller problem based on a random sample from the set of possible scenarios. Let $\xi^1, ..., \xi^N$ be an i.i.d. random sample of N realizations of the random vector ξ. Then the SAA problem can be defined as:

$$\min_{\mathbf{x}\in\mathcal{X}}\{\hat{g}_N(\mathbf{x}) = \frac{1}{N}\sum_{l=1}^{N} G(\mathbf{x},\xi^l)\} \tag{31}$$

If v^* and $\hat{v}_N$ represent the optimal values of the "true" and SAA problems respectively, it is well known that $\hat{v}_N$ is a valid upper statistical bound for v^* . Hence, the choice of large values of N will lead to better approximations of the true objective function. However, since the computational complexity of the SAA problem increases exponentially with the value of N, it is more efficient to select a smaller sample size N, and solve several SAA problems with i.i.d. samples.

Let M represent the number of SAA problems solved, and let $\hat{v}_N^m$ and $\hat{\mathbf{x}}_N^m$, $m = 1, \ldots, M$, denote the optimal objective value and solution of the mth replication, respectively. Once a feasible solution $\hat{\mathbf{x}}_N^m \in \mathcal{X}$ is obtained by solving the SAA problem, the objective value $g(\hat{\mathbf{x}}_N^m)$ need to be calculated or approximated by the unbiased estimator

$$\hat{g}_{N'}(\hat{\mathbf{x}}_N^m) = \frac{1}{N'} \sum_{l=1}^{N'} G(\hat{\mathbf{x}}_N^m, \xi^l) \tag{32}$$

where N' is typically larger than N, as the computational effort required to estimate the objective value for a given solution is generally less than that required to solve the SAA problem. One would also want to estimate the quality of the solution $\hat{\mathbf{x}}_N^m$. This can be done by computing an estimate of the optimality gap $v^* - g(\hat{\mathbf{x}}_N^m)$, where $g(\hat{\mathbf{x}}_N^m)$ can be estimated by (32), and v^* can be calculated exactly or approximated by

$$\bar{v}_N^M = \frac{1}{M} \sum_{m=1}^{M} \hat{v}_N^m \tag{33}$$

The sampling procedure can be terminated once the optimality gap estimate is sufficiently small or after performing all M replications, and the best solution among the SAA solutions can be selected using an appropriate criterion.

Effective implementation of the above sampling procedure requires that the SAA problems can be solved efficiently for relatively large values of the sample size N. This is especially suitable for (17)-(29), as it is relatively easy to evaluate the second stage objective function for given values of the $\mathbf{x}$ vector. Hence, for this problem we evaluate the value of a given portfolio exactly, rather than approximating it through sampling. A similar sampling based procedure is also implemented on a more generalized multiple stage R&D portfolio optimization model in [**?**].