October, 2009

Carbon capture and storage: Combining economic analysis with expert elicitations to inform climate policy

Erin D Baker, *University of Massachusetts - Amherst*
Haewon Chon, *University of Maryland - College Park*
Jeffrey M Keisler, *University of Massachusetts Boston*

Carbon Capture and Storage: Combining Economic Analysis with Expert Elicitations to Inform Climate Policy

Erin Baker*, Haewon Chon† and Jeffrey Keisler‡

December 10, 2008

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

†Haewon Chon, Joint Global Change Research Institute, University of Maryland, 8400 Baltimore Avenue, Suite 201, College Park,
MD 20740; hwchon@gmail.com

‡Jeffrey Keisler, College of Management, University of Massachusetts, Boston, MA 02125; jeff.keisler@umb.edu

Abstract

The relationship between R&D investments and technical change is inherently uncertain. In this paper we combine economics and decision analysis to incorporate the uncertainty of technical change into climate change policy analysis. We present the results of an expert elicitation on the prospects for technical change in carbon capture and storage. We find a significant amount of disagreement between experts, even over the most mature technology; and this disagreement is most pronounced in regards to cost estimates. We then use the results of the expert elicitations as inputs to the MiniCAM integrated assessment model, to derive probabilistic information about the impacts of R&D investments on the costs of emissions abatement. We conclude that we need to gather more information about the technical and societal potential for Carbon Storage; cost differences among the different *capture* technologies play a relatively smaller role.

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

Keywords: Climate change; Technology R&D; Uncertainty; Environmental policy

1 Introduction

In this paper we combine expert elicitations and economic modeling to analyze the potential for Research and Development (R&D) into Carbon Capture and Storage (CCS) to impact climate change. We concentrate on R&D investment directly, leaving the analysis of the government’s role in R&D investment to future research. We consider how R&D impacts technical change, and how technical change impacts the cost of reducing carbon emissions. We focus on technological improvements to *carbon capture* technology, while using some prior results in the literature on the feasibility and acceptability of long term carbon storage. We study the impacts of technical change on the entire abatement cost curve, which measures the costs of abatement, defined as a reduction in greenhouse gas emissions, at each level of abatement between zero and 100%.

This paper is part of a larger research project in which we are analyzing a number of climate change energy technologies including solar photovoltaics [3], nuclear power [2], electricity from biomass, wind and solar grid integration, liquid bio-fuels, and battery technologies. See [4] for an overview of the framework.

Crafting policies to address climate change presents us with a classic problem of dynamic decision making under uncertainty. We are faced with deep uncertainty about climate damages, in particular about how an extra ton of carbon emitted today will result in a stream of damages in the future. This is what economists call *marginal social damages*, and is crucial for determining climate policy. We are also uncertain about the future of technical change. CCS is particularly interesting because in many Integrated Assessment Models (IAMs), CCS plays a very large role in allowing the world to stabilize emissions in a reasonable period of time at a reasonable cost (See [12] and [27]). Yet, CCS at the scales imagined in the IAMs is far from an established technology. There is a great deal of uncertainty about our ability to build carbon capture plants at the costs assumed by many IAMs, and there is uncertainty about the technical and social viability of large scale storage.

The Role of the Marginal Abatement Cost Curve. 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. We have shown in past work [5] that it is particularly important to understand how new technologies will impact the Marginal Abatement Cost Curve (MAC). This is the curve that reflects the cost of reducing emissions by an additional ton. The MAC can be combined with uncertain marginal damages to analyze climate policy – both emissions policy and technology policy. The output of this paper is to show how different R&D investments in CCS technologies will lead to different probability distributions over MACs.

We note here that there are many ways to evaluate new technologies in terms of climate change. One prominent example is estimating the present value of welfare gain under the assumption of a given stabilization target. We are focussing, however, on the impacts to the MAC for several reasons. First, many abstract, analytical models represent technical change in terms of its impact on the MAC [9][19][21][22][29][33][41] or on the abatement cost curve directly [1][6][23][34][38]. Our work provides an empirical basis for the representation of technical change in this literature. Moreover, understanding the impacts on the MAC is important because, when combined with a marginal damages curve, it determines the optimal amount of abatement, and implicitly, the optimal amounts of different technologies to be used in the economy. Thus, we are uncertain about which stabilization target we will ultimately aim for, because we are uncertain about the marginal damages from climate change AND we are uncertain about which technologies will be available to mitigate climate change. 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. Thus, we will investigate here the impact of technical success in the defined CCS technologies on the MAC.

Organization of Paper and Flow of Data. In [4] we described a general framework for quantifying the uncertainty in climate change technology R&D programs and their associated impacts on emission reductions. In [3] we present an implementation of that framework, focusing on advanced solar technology; and here we focus on CCS. Figure 1 illustrates the flow of the data in this framework; the actions placed within the box are discussed

Figure 1: A schematic representation of the flow of data in the framework. The elements inside the box are explicitly discussed in this paper.

in this paper; the actions outside the box are applications of the outputs of this paper. The first step of the project, discussed in Section 2, is collecting probabilistic data on CCS technologies through expert elicitations. The products of the elicitations include explicit definitions of endpoints for each technology, and probabilities of achieving those endpoints for given funding trajectories. In Section 2.7, we compare our results with a similar study performed by the National Academy. In Section 3 we determine how the technologies would impact the abatement cost curve, if they achieve the defined endpoints. For this step we use MiniCAM, a technologically detailed Integrated Assessment Model, to determine the impact of each technology on the Marginal Abatement Cost Curve. In Section 3.4 we discuss the parameterization of each technology’s impact on the MAC. In Section 4, we combine the probabilities with the impacts on the MAC to derive multiple representations of the probabilistic impacts of R&D. As shown in Figure 1, these can then be combined with probability distributions of climate damages in technology policy models. We conclude with a discussion about further work that needs to be done in regards to gathering data about CCS in Section 5.

Richard Doctor	Argonne National Laboratory
Barry Hooper	Cooperative Research Centre for Greenhouse Gas Technologies
Wei Liu	Pacific Northwest National Lab
Gary Rochelle	University of Texas at Austin

Table 1: CCS Experts

2 Expert Elicitations

In this section we discuss the steps in the expert elicitation that we performed. These include the selection of particular technologies, the development of definitions of endpoints for the technologies, how we structured the probability assessments, and the results of those assessments. The output from the expert elicitations are specific definitions of success for each R&D project, and probabilities of success for each of those projects, given specific funding trajectories. We then go on to compare our results with the National Academy Study.

2.1 Identifying Experts

After a preliminary review of the literature, we contacted a set of engineering and economic experts who had familiarity with the range of CCS technologies. First, we met in person with experts who were especially knowledgeable in particular areas to structure the assessments as described below. We then identified 4 different experts with a range of areas of expertise (within CCS) to provide the probability judgments through completing written questionnaires. (While more experts would always be desirable, as we shall discuss with regard to future research, this small sample size is useful for obtaining preliminary estimates of mean probabilities.) We note that researchers may tend to self-select to work on technologies they believe hold promise, but those closest to a technology may also be most knowledgeable about its risks and the prospects for overcoming those risks. We balanced these considerations by working with a mix of CCS experts, by discussing with each expert potential biases and how to give answers as objective as possible, by obtaining rationales for assessments and by facilitating a round of feedback in which experts challenged each others' opinions. See Table 1 for a list of the experts who provided probabilities.

2.2 Technologies Considered

Our focus is on understanding how current investment in R&D has the potential to lower abatement costs 40 to 50 years in the future. Hence, we asked the initial set of experts to identify areas where there was potential for significant progress or even breakthroughs within this time frame. CCS technologies involve removing carbon during the production of energy so that it can be stored rather than released into the atmosphere. There are three main categories of CCS corresponding to three points in the process: Pre-combustion carbon capture, alternative combustion, and Post-combustion removal. Several key technologies within each area were identified, in part so that we could consider the potential return from a range of investment levels. These technologies (as well as a more comprehensive list of developing technologies) are described at length in, for example, [32],[26] and [30].

Pre-combustion capture works in conjunction with combined cycle power plants, such as IGCC (Integrated Gasification Combined Cycle) or NGCC (Natural Gas Combined Cycle) to remove CO₂ from syngas generated from fossil fuels or biomass, resulting in a stream of hydrogen to be used for combustion. Challenges are to make this process energy efficient and robust.

Technologies that remove CO₂ during combustion are more varied and involve changing the environment in which combustion occurs so that carbon is not released. We originally considered two promising directions. Chemical looping uses fine solid particles to carry oxygen to react with the fuel and then carry CO₂ away from the reaction without release into the air. This technology faces the challenge of finding particles that are both effective and durable at high enough temperatures. If these challenges can be successfully met, however, it is a very attractive technology, with much lower energy and non-energy demands. Using pure oxygen for combustion results in a flue gas that is almost entirely CO₂ and H₂O, thus requires almost no additional energy to separate out CO₂ [28]. Supercritical water oxidation (for which we modeled the impact of a successful technology, but did not assess probabilities) dissolves CO₂ as it forms, allowing it to be processed without being released as a gas.

Post-combustion CO₂ separation removes CO₂ from flue gases, through several possible means: membranes, solvents, stimulus or cryogenic methods. Its challenges involve finding materials that are effective, safe, and inexpensive to create and operate, even on existing power plants. This covers (at a high level, and not exhaustively) a wide range of capture technologies in various stages of R&D.

2.3 Definitions of Success

In order to meaningfully assess probabilities of success, the events constituting a successful endpoint must be defined unambiguously enough that one could say, after the fact, whether or not the event has occurred [44]. This endpoint is not meant to imply any absolute threshold between success and failure, but rather an operational definition for the purpose of assigning probabilities. Although it is theoretically possible to fix any arbitrary set of values as an endpoint, input from one or two CCS experts per technology was necessary to identify reasonable endpoints – endpoints that are not so ambitious as to be practically impossible, but that require some scientific advances beyond incremental improvements. Note that in order to translate the impact of technical success onto the entire MAC, we did *not* condition our elicitations on a specific carbon price. We discuss the ramifications of this in more detail in Section 2.7 below.

We define success across multiple dimensions of performance, to the extent possible in terms that will be easy for experts to relate to technology research activity and also practical to translate into parameters for economic analysis. This led to the use of differently defined technical hurdles for expert assessments on different technologies, even though we ultimately converted these costs to comparable terms to calculate the MAC curve. The definitions were as follows:

- Pre-combustion carbon capture:
 - *Assuming* that IGCC technology becomes a standard; and
 - *assuming* design is optimized to 90% capture;
 - parasitic energy loss (percent of energy produced which is devoted to capture rather than production of electricity) $\leq 10\%$; and
 - incremental capital cost for IGCC $\leq 10\%$.
- Chemical looping:
 - Operation at 1200 degrees K (so that combustion will be efficient); and
 - cost of energy of 0.05cents/kWh or less. Note, the cost of energy for Chemical looping includes capital, fuel, maintenance and operation; these are inter-related and will require either development of durable

absorbent particles or a design that allows for easy replacement of absorbent particles.

- meeting environmental regulations for other pollutants.
- Post-combustion CO₂ separation: a process that achieves
 - availability of 90%, (i.e., maintenance of 2 weeks per year, comparable to existing plants);
 - derating (amount of additional energy compared to a no-capture plant) of no more than 30%, assuming a design optimized for a capture rate of 90%;
 - cost per ton of CO₂ avoided of \$25 per ton or less;
 - on at least 50% of available coal (that is, on coal other than lower-quality lignite).

We also defined a technical endpoint for Super-critical water oxidation, but due to the specificity of this technology, we were not able to find sufficient experts to assess it. We include the definition here for completeness:

- Super-critical water oxidation:
 - Fuel processing temperature < 850 degrees K;
 - Less aggressive target:
 - * combustion temperature ≥ 1400 degrees K;
 - * and COE $\leq \$0.05/\text{kWh}$,
 - More aggressive target:
 - * combustion temperature ≥ 1800 degrees K; and
 - * COE $\leq \$0.05/\text{kWh}$;
 - Additionally, aquifers must have a usable lifetime ≥ 40 years (before they are saturated)
 - 100% capture rate.

We defined two targets – less aggressive and more aggressive – to account for a trade-off between efficiency and plant lifetime. At higher temperatures, the process would be more efficient, but more corrosive, thus reducing plant lifetime. Hence, we considered a lower temperature target which would be easier to achieve, but less likely

to lead to a cost advantage even if it does succeed; and a higher temperature which would be harder to achieve, but more likely to lead to the low cost if it succeeds. Of particular note with this technology is that it is defined to achieve a 100% capture rate, if successful. In our analysis below we illustrate the relevance of this.

To put these definitions in perspective, consider the estimates presented in IPCC 2005 [32]. They report current energy requirements and incremental capital costs for current IGCC plants of 14-25% and 19-66% respectively. Our targets for pre-combustion represent a decrease of 29% and 47% below the low estimates. This compares with the IPCC prediction that, with sustained R&D, costs will reduce by 20-30% in the next 10 years. Thus, the targets defined above appear to be in line with the IPCC report. For post-combustion, the IPCC reports energy requirements and cost per ton of CO₂ avoided for current pulverized coal plants of 30% and \$30 - \$70, respectively. Our targets are roughly equivalent in terms of energy requirement, and represent about a 17% reduction in costs. Thus, our target does not appear overly ambitious in comparison with the IPCC report. There are no good comparisons for Chemical Looping or Super-critical Water Oxidation.

2.4 Probability Elicitations

We now describe the elicitation of probability assessments for the technologies in this study, describing the process of structuring these assessments and conducting surveys to obtain judgments from multiple experts. Our method for assessing prospects of long-term R&D under varying funding scenarios is based on insights from the standard Decision Analysis literature on obtaining probability judgments from experts in ways that avoid biases due to overconfidence, motivation, anchoring, etc. [35][44][45]. In the following sections we present the raw results of these surveys and discuss how multiple probabilities may be combined.

As mentioned above, we worked with an initial set of experts to structure event trees describing which combinations of events must occur to achieve success as defined above. We developed detailed influence diagrams [25], and then distilled these into event trees to on which to base our probability assessments. Before administering the final assessments, we conducted practice assessments where practical with the experts who developed the structure. We also previewed the final questionnaire with one of the second set of experts. From this preview, we then modified some of the questions to improve their formulation before the final elicitation. The structure of the event trees varied across the technologies. For Pre-combustion, we simply assessed the likelihood of achieving

Figure 2: Structure of probability tree for Post-Combustion

Figure 3: Structure of probability tree for Chemical Looping

the defined levels of energy loss and capital cost, under the assumption that IGCC was widely used and that the process was optimized for 90% capture. Thus, it had two hurdles to overcome in series. Where multiple paths are available to achieve success (e.g., Post-combustion separation, Figure 2), we established that the probability of success on each of the pathways was independent of success on other pathways. Then, with probabilities for each pathway, the total probability of success is 100 percent minus the product of the probability of failure for each pathway. Note that in this tree, the probability for each pathway takes into account how likely that pathway is to clear all the relevant hurdles (cost, derating, etc.).

Chemical looping (Figure 3), has two possible paths to success that share a first hurdle (high temperature performance) followed by cost and environmental hurdles for two variations of the technology: with and without dynamic replacement of particles. In this case, the latter probabilities can be different depending on which variation of the technology is considered. After calculating overall probability of success for each technology from the probabilities for overcoming the hurdles, we compared these implied probabilities against overall holistic estimates of a technology’s probability of success as a consistency check to offset potential decomposition bias [43] in which adding details lowers overall probability estimates.

In addition to defining endpoints for each technology, experts defined possible U.S. government funding trajectories. These provide the basis for estimating probability of success. Specifically, we formulated a “baseline” funding trajectory which would give a fair chance to show whether the technology could succeed, i.e., if the technology is really good, the funded research would most likely demonstrate that; a high funding trajectory, beyond which it would be hard to fund good research; and a low funding trajectory, which would at least allow research to continue. We did not explicitly consider a zero-funding trajectory. Funding is defined in terms of duration, amount, and source. Because these trajectories take into consideration at the outset the range of research possibilities, each technology area had a different set of trajectories. As with the endpoints, these funding trajectories are not meant as forecasts or as definitive statements about what is required, but rather function as working assumptions in order to facilitate the estimation of probabilities. These trajectories are given with the assessment results. For Pre-combustion we considered 10-year funding of either \$5, \$20, or \$50 Million dollars per year. For Chemical looping, we considered 10-year funding, with a lower amount in the first five years, and higher amount for the 2nd five years. The low trajectory was \$0.5 Million/year followed by \$5 Million/year; the

baseline trajectory was \$1M followed by \$10M; and the high funding trajectory was \$5M followed by \$10M. For Post-combustion we considered a 15 year funding trajectory. The baseline trajectory was defined as \$15M/year for the first year, ramping up at an even rate to \$30M/year in the last year. The low trajectory was simply \$5M/year; and the high trajectory \$50M/year.

We then sent out a survey to the second set of experts (listed in Table 1). These experts were all comfortable working with probabilities. In addition, before they answered these questions, they read a primer we prepared on probability assessment which described potential biases that can arise and how to avoid them. They provided written rationales for the estimates.

2.5 Assessment Data

Tables 2, 3, and 4 below summarize the set of assessments, giving each expert’s probability of each hurdle being achieved for each technology, at three funding levels each. Expert 4 only felt comfortable assessing Post-combustion. Expert 2 only assessed one question on Chemical looping.

Funding Trajectory	<i>\$5M per year</i>			<i>\$20M per year</i>			<i>\$50M per year</i>		
	Ex 1	Ex 2	Ex 3	Ex 1	Ex 2	Ex 3	Ex 1	Ex 2	Ex 3
Parasitic energy loss < 10%	5%	50%	2%	15%	70%	10%	20%	90%	25%
Increase capital cost < 10%	66%	10%	2%	70%	30%	7%	90%	50%	15%
Total Probability	3%	5%	0%	11%	21%	1%	18%	45%	4%

Table 2: Summary of Assessment Results for Pre-combustion

Funding Trajectory	<i>\$0.5M to \$5M</i>			<i>\$1M to \$10M</i>			<i>\$5M to \$10M</i>		
	Ex 1	Ex 2	Ex 3	Ex 1	Ex 2	Ex 3	Ex 1	Ex 2	Ex 3
Operate at 1200K or higher	50%	N/A	1%	85%	95%	2%	95%	N/A	10%
5 c/kwh with dynamic replacement	50%	N/A	1%	85%	N/A	2%	95%	N/A	5%
5 c/kwh w/o dynamic replacement	40%	N/A	1%	75%	N/A	2%	85%	N/A	5%
Meets env.regs wth replacement	50%	N/A	1%	75%	N/A	2%	90%	N/A	5%
Meets env.regs w/o replacement	35%	N/A	1%	50%	N/A	2%	60%	N/A	5%
Total Probability	16%	N/A	0%	59%	N/A	0%	84%	N/A	0%

Table 3: Summary of Assessment Results for Chemical Looping. Please note that probabilities in the bottom row are rounded,

The most striking result is the significant disagreement between the experts on the probability of achieving

Funding Trajectory	\$5M per year for 15 years				\$15M per year ramping up to \$30M (15 years total)				\$50M for 15 years			
	Ex 1	Ex 2	Ex 3	Ex 4	Ex 1	Ex 2	Ex 3	Ex 4	Ex 1	Ex 2	Ex 3	Ex 4
Molten or alternative solvents	33%	25%	0%	1%	75%	50%	2%	2%	90%	75%	5%	4%
Cryogenic methods	75%	50%	5%	5%	90%	75%	7%	10%	95%	95%	10%	15%
Stimulus methods	0%	10%	2%	5%	0%	20%	7%	10%	0%	50%	15%	15%
Ammonia/amine membranes	33%	95%	5%	30%	75%	95%	7%	50%	90%	95%	10%	70%
Total Probability	89%	98%	12%	37%	99%	100%	21%	60%	100%	100%	35%	79%

Table 4: Summary of Assessment Results for Post-combustion. Please note that probabilities in the bottom row are rounded.

the given endpoints. Experts 1 and 2 appear to be optimists; while experts 3 and 4 are pessimists.

Figure 4 presents a graphical representation of the results. It shows how the probability of achieving each endpoint is related to the net present value of the R&D investment trajectory (assuming a 5% real discount rate). We have averaged the two optimists’ results, and separately, the two pessimists’ results. One question of interest here is whether the disagreement is fundamental – reflecting different beliefs about the overall viability of the defined endpoints; or whether this is a disagreement about how much funding is required to achieve these endpoints. In this case, (in contrast to what we found for solar [3]) the comments of the experts indicate that the disagreement is fundamental. Regarding the sub-technology of Post-combustion, *Molten or alternative solvents*, one expert, an optimist, says “The definition of success is close to currently available commercial systems. Workable technology for molten carbonate fuel cells;” while another, a pessimist, says “\$25/ton will be practically impossible with the best of technologies. Even with perfect membranes, the cost of compression will be high.” Disagreements such as these are found throughout the surveys. We circulated these comments among all the experts, but found very little change in their assessed probabilities.

As the second quote above illustrates, one of the fundamental disagreements is whether the specified cost goals can be met. In [3] we found this to be the case in our expert elicitations about solar. In fact, disagreements over the probability of achieving cost goals appears to be emerging as a theme in all the expert elicitations we are doing. This reflects at least two issues. The first is that many of our experts are scientific experts, familiar with the relevant scientific properties, but perhaps less familiar with all the factors that determine ultimate cost. The second issue is that trying to predict costs is very complex, at least as complex as trying to predict scientific breakthroughs. The ultimate cost of a technology will depend on its design and the materials used to make it;

Figure 4: Assessment Results

but also on the specific processes employed to produce it, and on returns to scale. See Rubin et. al. [42] for a discussion of learning curves and CCS; and Riahi [40] for a discussion of returns to scale. This presents a very challenging problem in eliciting probabilities over future technologies. If a cost goal is *not* specified, then many of the probabilities of success would be very near 1 – scientists can achieve a wide range of results if cost is no object. Moreover, some scientific breakthroughs are crucial to determining the cost. For example, as we mentioned above, achieving success at a higher temperature leads to higher efficiency and therefore impacts cost. Yet, scientific experts may not be the appropriate ones to assess the future economic interactions and returns to scale that will determine ultimate cost.

How to approach this is an open question. See Nemet and Baker [37] for one framework that combines expert elicitations on scientific breakthroughs (not conditioned on a carbon price or other economic variables) with economic cost models that are conditional on carbon price and other economic variables that impact demand, and therefore returns to scale. Alternatively, economic or business specialists could be elicited along with scientists and engineers.

The elicitation results also shed light on whether returns to R&D investment are increasing or decreasing in scale. Most of the responses show decreasing returns over the range of investments we have specified. We note, however, that we did not explicitly ask for the probability of achieving the endpoints in the absence of government R&D funding. In the case of Chemical looping, it appears that the returns are S-shaped over this range, with

strongly increasing returns for the early investments, tapering somewhere around \$30M. On the other hand, Pre-combustion seems to be showing some increasing returns to scale. The optimists show decreasing returns to scale, but just barely. The pessimists show increasing returns to scale over the range that we assessed. The impression is that it would be beneficial to assess a higher investment trajectory for this technology.

The wide range of opinion within the research community is itself notable, and suggests that the connection from research to development is not yet well understood. Future research could explore further the reasons for disagreement. For example, panels could be convened to discuss a range of endpoints.

2.6 Combining Expert Judgments

For modeling purposes, we can compute returns to R&D for the technologies assuming various combinations of the elicited probabilities. Most simple is to calculate an overall probability of success for each technology, for each expert based on that expert’s expressed probabilities regarding each hurdle. This can be used for sensitivity analysis, but gives quite wide ranges. In this paper we will focus on two results for illustrative purposes – the simple average of all experts; and the log-odds average. Given the wide range of expert responses, it is best to interpret these averages cautiously [31]. Alternatively, we could get different results from the same data, e.g., by averaging probabilities at the hurdle level rather than the technology level, or giving different weights to different experts [15]. The simple average presented is responsive to all judgments but not prone to large swings based on a single opinion [45].

The log-odds ratio is the natural log of the odds of success, derived by dividing the probability of success by the probability of failure. We calculate the average of the log-odds ratio for each expert, then back out the probability implied by this. This method is much more sensitive to very large and very small probabilities. Figure 5 compares the simple average with the log-odds ratio average for each of the technologies. For Post-combustion, the log-odds ratio average is higher than the simple average, since the optimists have very high probabilities, approaching 100%. For the other two technologies the log-odds ratio average is lower than the simple average, because the pessimists have very low probabilities, approaching 0%. The difference between the two averages is most striking for Chemical looping, where the simple average of the probabilities of success given the high funding trajectory is 42% while the log-odds ratio average is 4%.

Figure 5: Combined Expert Judgements

2.7 National Academy Study

In 2007 the National Academy released a report on the 2nd phase of a study on prospective evaluation of applied energy R&D at DOE [36]. In this report they recommend the use of expert elicitations in order to perform uncertainty analysis over prospective benefits from R&D programs. In this 2nd phase the researchers performed a case study on the DOE’s carbon sequestration program, using a panel of 14 experts. In this section we compare the results of the case study with our results.

2.7.1 Carbon Capture Costs

The NAS study focused on implementation of CCS technologies by 2025. They compare a no-DOE funding scenario with the DOE program as currently projected to be funded. The DOE projects spending an average of about \$220M a year over a 20 year period starting in 2001, on the CCS program. This program, however, includes six research streams: CO₂ capture; carbon storage; monitoring, mitigation and verification; breakthrough concepts; non-CO₂ greenhouse gases; and infrastructure development. The first research stream is most similar to the technologies we developed. The research stream on “breakthrough technologies” also considers CO₂ capture technologies and so is relevant.

While we define success endpoints in terms of specific technology, the NAS considered the overall success in terms of increase in cost of electricity (COE). Four possible choices of increases in COE were given to the expert

panel: less than 10%; between 10-20%; between 20-30%; or greater than 30%. The mean increase in COE from the panel assessments was about 15-18% with the DOE program and about 16-20% without it, depending on the level of carbon tax [36].¹ In comparison, the technologies we considered have COE increases for coal-fired plants of about 7.5% for Chemical looping, 20% for Pre-combustion, and 35% for Post-combustion².

Figure 6 compares the average probabilities in the NAS study with ours. The left panel compares the NAS assessments assuming no DOE funding, with our assessments assuming a low funding trajectory for each technology, a total of about \$15M per year. The right panel compares the NAS assessments assuming DOE funding (\$220M over six categories), with our assessment assuming a high funding trajectory for each technology, a total of about \$110M per year. We have used the mid-points of each range for the NAS figures; that is, we have represented “less than 10%” by 5%, etc. The figures show the probability of achieving the additional cost to COE on the horizontal axis, or better. So, for example, in the left panel with no or low funding, the NAS has a probability of achieving an additional cost of 15% or less of about 55%; our experts say about 7.5%.

Despite the fact that our funding amounts are higher and our implementation date is later, the average NAS assessments are more optimistic than our experts on average in almost all cases. This may reflect that we assessed specific technologies to reach each of the cost goals, whereas the NAS study assessed aggregate increases in COE, leaving the choice of technology flexible. On the one hand, we would expect their probabilities to be somewhat higher since they are not limited to our specific technologies. On the other hand, it is possible that the NAS experts were not thinking about the specific hurdles that need to be overcome, and so were over-optimistic. It also may reflect the different ways the questions were worded. In particular, we defined success in Post-combustion technology in terms of achieving a cost of \$25/ton of avoided CO₂. This struck some of our experts as a very ambitious goal, with one expert claiming “practically impossible”; and even our optimistic experts did not assign a probability of 100%. However, this cost of avoided CO₂ is equivalent to about a 35% increase in COE. In the NAS study, the *worst* the technology could do was an increase of 30% or greater. Therefore, this level of increase is essentially in the “failure” bin. It may be that the NAS experts were not carefully translating back and forth

¹The carbon tax is assumed to impact the private sector incentive to innovate; but does not impact the increase in COE directly.

²These estimates of COE only include additional cost of capture and compression, but exclude the cost of geologic storage. They are presented only for comparison purposes. The actual COE may vary by fuel type, sequestration method and distance, and changes in carbon price and demand for fuel. The estimates presented here assumes zero carbon price and a baseline coal price of \$1.7/GJ of primary energy. The detailed assumptions and calculations are presented in the following section.

Figure 6: The cumulative probability over the additional COE needed to achieve CCS.

between an overall increase in COE and a cost of CO_2 avoided. Our experts were more optimistic (on average) in only one case: they put a high probability of success on Chemical looping under the high funding trajectory.

Finally, a key difference between the two studies is that the NAS results are conditional on specific levels of a carbon tax (\$100 and \$300/tC); while our elicitations are specifically *not* conditional on a carbon price. Whether or not to condition on a carbon price is a tricky methodological question. The NAS included it because they were very concerned about capturing the effects on the technology of private sector investments; and private sector investments are strongly dependent on the carbon price. We did not include it, because our goal is to determine the (probabilistic) impact of R&D investments on the entire MAC. The MAC, in turn, impacts the optimal carbon price. Thus, we did not condition on a carbon price in order to avoid circularity. A clear direction for future work is to simultaneously capture the effects of incentives on private sector innovation and the impact of innovation on the future MAC.

2.7.2 Viability of Long Term Storage

The DOE Carbon Sequestration and Technology Roadmap [17] lists a number of challenges to the widespread implementation of CCS, including permanence; monitoring, mitigation and verification; permitting and liability; and public acceptance. The NAS study explicitly considered public opposition based on the risk of sequestration; regulatory issues; and physical siting requirements. We did not address these issues in our elicitations. The NAS study addressed it very briefly. 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.” They note that there was wide disagreement on this, but that the lowest assessment was above 0.5. We will use these estimates in our analysis below.

3 Impact on the MAC

The expert elicitations provide us with definitions of success for improvements in CCS technology. This section describes the applied analysis of the impacts that success, as defined in the elicitations, might have on the costs of CO₂ abatement. We start by converting the definitions of success given above into a consistent set of parameters for use in MiniCAM. We then go on to derive a Marginal Abatement Cost curve for each definition of success.

3.1 From Definitions of Success to MiniCAM Parameters

In this section we present the conversion metric we used in order to obtain the parameters for modeling with MiniCAM. In MiniCAM, CCS technologies are characterized by three technical and economic parameters: capture rate, parasitic energy loss, and non-energy cost. Capture rate is defined as the percentage of carbon captured in the CCS process per ton of carbon contained in the fuel input. Parasitic energy loss is the proportion of reduction in electricity output due to capture process. Non-energy cost is defined as other costs associated with establishing and maintaining the carbon capture component of the power plant, but not the cost of transport and storage of captured CO₂. The cost of transport and storage of CO₂ is not subject to technical change in our analysis. We used a constant cost of \$58/tC^{3 4} obtained from the Climate Change Science Program [13] throughout this paper.⁵

For Pre-combustion, a 10% parasitic energy loss and 90% capture rates were obtained directly from the expert elicitation. (All three technologies considered share the same capture rate of 90%.) Non-energy cost needs to include both capital cost and O&M cost. The success endpoint also specified a 10% increase in the capital cost over the plant without capture. Our baseline capital cost for a coal-fired IGCC power plant obtained from CCTP

³Except in South Korea and Japan, where the storage cost is assumed to be 5 times higher due to scarcity of land.

⁴This figure is higher than the commonly used \$10/tCO₂ figure [16], which is equivalent to \$37/tC (multiplying by the ratio of molecular weight of CO₂ (44) and of carbon (12)); and to \$41/tC in 2005 constant dollars.

⁵Throughout this paper we use 2005 constant dollars and metric tons, unless otherwise indicated.

[14] is \$2.7 cents/kWh. A 10% increase in the capital cost corresponds to 0.27 cents/kWh. For variable non-energy cost not specified in the elicitations, such as O&M cost, we used the baseline figure from CCTP of 0.16 cents/kWh. Summing these two yields total non-energy cost of 0.43 cents/kWh.

For the Post-combustion technology, experts denoted the energy requirement of capture in terms of derating. A 30% derating is equivalent to 23% parasitic energy loss.⁶ The elicitation also specified a total capture cost (including energy and non-energy costs) of \$25⁷ per ton of CO₂. Assuming the carbon content of coal of 27.3 kgC per GJ of primary energy and 50% thermal efficiency (LHV), the total capture cost is equivalent to 1.57 cents/kWh. We used a baseline power plant electricity cost of 4.5 cents/kWh from CCTP [14] to convert 23% parasitic energy loss into 1.05 cents/kWh. By subtracting the cost of parasitic energy loss from the total capture cost, we obtain non-energy cost of 0.52 cents/kWh.

For chemical looping, the final cost of energy is given at 5.0 cents/kWh in 2006\$, or 4.8 cents/kWh in 2005\$. This is merely 0.28 cents/kWh higher than our reference coal-fired power plant in the model (in the year 2050). Since the technology is still in the experimental stage, it is difficult to find a reliable ratio between the energy cost and the non-energy cost. We adopted the ratio from the MiniCAM's baseline CCS component [14]; 57% for non-energy cost and 43% for energy cost.⁸ 57% of 0.28 cents/kWh is 0.16 cents/kWh for additional non-energy cost. The remaining energy cost portion of 0.12 cents/kWh is divided by the baseline electricity cost of 4.5 cents/kWh to obtain 2.6% in parasitic energy loss. Since the chemical looping technology is not an add-on component of non-CCS power plant, but rather a different type of power plant, we can think of it as a new type of power plant with combined non-energy cost of 3.5 cents/kWh and 49% efficiency (LHV).

CCS technology can also be applied to other fuel sources, namely oil and gas.⁹ We did not explicitly define technology endpoints for these other fuel sources. Rather, we assumed that they would show improvements proportional to the improvements in coal. Specifically, we adopt the ratios between coal and other fuel sources for non-energy cost and parasitic energy loss from the Climate Change Technology Program [14]. These ratios are

⁶Derating measures the increase in primary energy input to obtain fixed final energy output, while parasitic energy loss measures the reduction in final energy output given fixed primary energy input.

⁷in 2006 constant dollars, when the questionnaire was developed

⁸This ratio between energy cost and non-energy cost prevents the CCS powerplant being superior to the non-CCS powerplant. While this ratio is somewhat arbitrary, due to the small difference in the cost of electricity, it is unlikely that any small change in this ratio would make a significant difference.

⁹CCS is also applicable to biomass, cement production, and other industrial processes. The scope of this paper only includes CCS for fossil fuel fired powerplant, but biomass CCS will be addressed in a future paper on biomass energy technologies.

then applied to the success endpoints derived for coal CCS technology to obtain parameters for oil and natural gas. These success endpoints are assumed to be widely applicable in the market by 2050. The energy requirement and non-energy cost for the three technologies and three fuel sources are shown in Table 5.

Technology		Post-Combustion			Pre-Combustion			Pre-Combustion w/ High-Capture			Chemical Looping		
Fuel		Coal	Gas	Oil	Coal	Gas	Oil	Coal	Gas	Oil	Coal	Gas	Oil
Parasitic Energy Loss	%	23%	20%	26%	10%	8.8%	11%	10%	8.8%	11%	2.6%	2.3%	3.0%
Non-Energy Cost	2005\$/kWh	0.52	0.57	0.74	0.43	0.47	0.61	0.43	0.47	0.61	0.16	0.17	0.23
Capture Rate	%	90%	90%	90%	90%	90%	90%	98%	98%	98%	90%	90%	90%

Table 5: Summary of Model Parameters for the year 2050

3.2 Methods and Assumptions

For this study, we derive MAC curves for the year 2050 under different assumptions about technological pathways. The analysis was conducted using the MiniCAM integrated assessment model. MiniCAM is a global IAM 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. MiniCAM explicitly represents a range of electricity-generating technologies including various generations of nuclear power, solar and wind power, electricity from biomass, as well as multiple fossil generating technologies, characterized by different fuel types and generation methods. Technology characteristics in MiniCAM are inputs to the model; the model does not include learning curves or other approaches to induced technical change. See Brenkert et al. [10] and Edmonds et al. [20] for more discussion of the model. Assumptions for technologies other than fossil CCS are based on the version of MiniCAM used in the Climate Change Science Program (CCSP) MiniCAM reference scenario [13].

The objective of the analysis was to develop MAC curves under specific assumptions about the characteristics of CCS technologies at a particular time in the future, in this case 2050. These curves relate levels of emissions reduction to carbon prices, thus they approximate the marginal cost of emissions reductions. A range of carbon price paths were created leading up to 2050. In each path, the carbon price increases over time at a constant

rate of 5%.¹⁰ In order to approximate MACs, the equilibrium emission abatement level in 2050 from the range of price paths are plotted on the horizontal axis against the carbon price in 2050 on the vertical axis (See Fig 8).¹¹ See [46] for another example of this kind of analysis. The relationship between abatement and the carbon price resulting from this analysis represents a marginal abatement cost function at a particular point in time, and based on the particular assumptions about capital deployment in previous time periods. A similar approach was used for analyzing the impact of cost reduction in photovoltaic cells [3].

The version of MiniCAM used in this analysis represents electricity produced by CCS technologies as having a constant unit cost. This cost is derived by summing the energy cost that varies with fuel cost, the carbon price, and the constant non-energy cost. Equilibrium output is derived by a logit formulation in the energy market, which captures regional and other factors that lead to heterogeneity in costs across applications [11]. For this reason, even if a technology such as CCS results in the lowest average cost of electricity due to high carbon prices, it will not capture the entire electricity market. This aspect of the model is quite distinct from the “winner-take-all” method used in some other analyses, including the aforementioned NAS study [36]. The logit formulation we use, although abstract, partially captures the issues of geographically varying CCS electricity costs, due to quality of fuel, access to fuel sources and sequestration site, etc. The same CCS technology would cost more in a location where coal mines and carbon sequestration sites are far away from where the electricity is demanded; therefore, even a CCS technology that has the lowest unit cost in the electricity market in general may still not be deployed in some areas. On the other hand, we have made the simplifying assumption that only the CCS technology with the lowest cost will be implemented: if all three technologies are successful, then chemical looping will be used; if both pre- and post-combustion are successful, only pre-combustion will be used.

Unlike some other low-carbon energy sources, such as wind and solar, CCS electricity is subject to a carbon tax. The CCS technologies considered do not capture 100% of the carbon in the fuel source. Thus, electricity produced by CCS technologies will still be subjected to carbon tax. According to the Global Energy Technology Strategy Program (GTSP) report, at high carbon prices, even a 10% release of CO₂ could have a profound impact on the selection of generation technology [18]. In fact, we do observe declining deployment of CCS at very high carbon prices associated with high levels of abatement from the modeling results (such as 90% abatement level). This

¹⁰This is a standard resource economics approach to resource extraction, see [24] and [39].

¹¹We define the MAC to be non-negative, and therefore do not show negative carbon prices.

Figure 7: Trajectories of parasitic energy loss and additional non-energy cost for the Coal-fired powerplant CCS technologies

effect is more pronounced for CCS technologies with higher capture energy requirement, such as Post-combustion technology. As more primary energy is required to produce a unit of electricity, more fuel is used, and thus more carbon is vented uncaptured. This effect dynamically changes the equilibrium level of deployment according to the carbon price and technology characteristics.

The treatment of technical change leading up to 2050 is also important. Since the lifetimes of power plants are considerably longer than the 15 year timesteps used in MiniCAM, power plants built in 2020 or 2035 would still affect the emissions reductions in 2050. Hence, in this analysis, instead of assuming an instantaneous advancement, the parameters of the CCS technologies are assumed to decline over time up to 2050. Figure 7 shows the trajectories of capture energy requirement and non-energy cost for coal-fired power plants with CCS. Natural gas CCS and oil CCS follow the same decline trend, conserving the ratio given in Table 5. Starting in 2020, Post-combustion energy requirement and non-energy cost are assumed to decline at a rate of 1% per year to meet the target parameters in 2050. Pre-combustion and chemical-looping technologies are assumed to share the same starting point in 2020, but the rate of decline is raised to meet the corresponding target in 2050.

Figure 8: MAC curves under different technology assumptions. The left and right panels show the MAC for abatement between 0%–50% and 50%–90% for emphasis.

3.3 MiniCAM Results

In this section we present the marginal abatement cost curves in 2050 generated by MiniCAM using the costs implied by the definitions of success in Table 5, and interpret the implications to CO₂ abatement.¹² Figure 8 shows the global economy-wide MAC curves generated based on the scenarios with each of the three CCS technologies introduced and on a reference scenario with the CCS component switched off. For purposes of comparison, we have also included an additional technology that has the parameters of Pre-combustion, but has a capture rate of 98%. We assumed that the CCS technologies are perfect substitutes. In joint success cases where more than two CCS technologies succeeded, we assumed only the lowest cost technology to be available.

Several observations of the results bear note. First, advances in CCS show minimal impact on emissions reductions in the absence of a carbon price. Such results are quite different from other alternative energy sources such as PV [3]. While low cost PVs may be deployed in geographically well-suited areas solely because they are the least expensive option in the market, it is unlikely for CCS, since even with large reductions in the cost, the final cost of electricity will always be higher than the cost of producing energy with the same technology without a CCS component. Other than small scale carbon capture to supply CO₂ for industrial uses [18], no other advantage would make CCS electricity less expensive than electricity from the same power plant with the

¹²All calculations are in terms of tons of carbon and 2005 constant prices.

CCS component turned off. Thus, CCS technologies will only make a visible impact on the MAC when the carbon price is positive.

Second, the absolute level of reductions in the MAC associated with the introduction of CCS technologies increases as the abatement level goes up. CCS technologies take a larger share of the electricity market as higher levels of abatement are required and the carbon price goes up. The benefit gained from reducing the cost of CCS electricity becomes larger as more CCS power plants are deployed. In fact, we observe the MAC curves with advanced CCS diverging from the reference MAC as the abatement level increases.

Third, the magnitude of the shift from the No CCS curve to the Post-combustion curve is larger than the shift between any of the two different technologies, beyond 40% abatement. At medium to high levels of abatement, incremental reductions in cost appear to be less important than the effect of simply having CCS technology in the market. At these levels of abatement scenarios without CCS start to run out of low-cost abatement options—such as fuel-switching to natural gas—and must start deploying higher cost options—such as further deployment of PVs in areas with less sunshine. Simply having CCS technology as an option, even at a high cost, relaxes such constraints.

Fourth, the availability of a very high capture rate has a large impact at very high levels of abatement. Note that the MAC for the high-capture technology veers off from the other MACs at about 80% abatement. At that point, the high capture rate (which allows this technology to avoid the carbon tax almost completely) outweighs the lower costs in the Chemical looping technology. If it is likely that very high abatement levels are to be warranted, it might make sense to focus some attention on technologies that have very high capture rates. More to the point, the ability to engineer very high capture rates at reasonable costs provides an *option value*. Hence, attention should be paid to any systematic differences between the technologies with respect to the ability to engineer high capture rates.

3.4 Parameterizing the Impact on the MAC

In this section we parameterize the impact of CCS on the MAC. We use the data generated above to estimate a smooth relationship between technical change and the impacts on the MAC. We do this for three reasons. First, it provides a single metric on which to do analysis, leading to insights about the cost effectiveness of different

types of R&D, for example. In Section 4 below we perform some initial analyses like this. Second, many top-down economic analyses represent technical change in terms of a shift or pivot to the MAC; thus our estimate provides an empirical bases for such analyses. Third, as illustrated by Figure 1, this work is intended to inform an overall portfolio analysis under uncertainty. For this we need to know how the entire MAC is affected; and we want it to be portable to a wide number of representations of the MAC. The first characteristic means that the impact on the MAC is the right level of detail. The 2nd characteristic argues for parameterizing the impact using a small number of parameters.

We observe from Figure 8 that CCS appears to pivot the MAC curve down. Thus, we parameterize the impact on the MAC through a pivot parameter. This parameter satisfies the above criteria in a way that the area under the curve, for example, would not. In terms of determining the optimal abatement, it is very important whether the new MAC diverges from the old at low levels of abatement or at high levels; a measure of area would not sufficiently indicate this.

Specifically, we let

$$\widetilde{MAC}(\mu; \alpha) = (1 - \alpha) MAC(\mu) \quad (1)$$

where the tilda represents the MAC after technical change parameterized by α . We estimate the parameter α by averaging over the pivot at different percentiles. We average the pivot for 3%, 5%, 7%, and 10% abatement. Then we average that with the pivot for 10%, 20%, ..., 90% abatement. We use this method rather than minimizing the square errors between the estimated curve and the empirical curve for two reasons. The first reason is that the MAC is much higher at high levels of abatement, and so minimizing the square errors will put too much weight on high levels of abatement. Second, the model becomes less stable at very high levels of abatement, and thus we have less confidence over that range. Averaging the pivots puts an even weighting on all levels of abatement. Table 6 shows the values of alpha for each of the three main technologies. Figure 9 compares the empirical MACs estimated using MiniCAM with the MACs estimated using the values of α in Table 6.

Technology	Chemical Looping	Pre-combustion	Post-combustion
Alpha	0.33	0.27	0.21

Table 6: Pivot parameters for each technology

In this section, we have estimated the impact of technical successes on the MAC curve. In the next section,

Figure 9: Comparison of Estimated MACs with actual MACs

we combine the assessed probabilities of achieving such success for the given funding levels with the calculated impact for a successful technology in order to analyze the net potential impact of R&D funding.

4 Data Analysis

In this section we combine the probabilities provided by the experts in Section 2.3 with information on the impacts on the MAC from Section 3. We want to represent the probabilistic relationship between R&D investments and technical change, where technical change is represented by the impact on the MAC, α . In order to translate our data into information about the returns to R&D, we need to hypothesize funding orders – rules determining which project will get funded first, second, etc. Using these funding orders, we present information on the marginal impact of additional R&D investment in two different ways – on the expected value of α and on the probability of success.

The most typical funding order would be to fund projects in the order of the expected impact per dollar invested. This is a good heuristic that is widely used in industry, although it clearly does not always result in

the optimal portfolio. One major weakness of this heuristic is that it ignores risk issues completely, only focusing on expected return. Therefore, it may be useful to consider alternate funding orders. We also consider a “high risk” and “low risk” funding order, since previous work has shown that the optimal riskiness of the projects may vary [1]. The “low risk” order is determined by funding projects in order of the probability of success per dollar invested. This is what an extremely risk averse decision maker might do. The high risk order is determined by funding projects in order of the potential impact per dollar invested (ignoring the probability of success). This corresponds to the behavior of an extremely risk seeking decision maker. For the Main and Low Risk funding orders, we have used both averaged probabilities and log-odds averaged probabilities. Table 7 shows the order in which projects would be funded for the five funding orders.

Name	Description	1	2	3	4	5	6	7	8	9
Main	Expected $\alpha/\$$, using average probabilities	Chem Loop, Med	Chem Loop, High	Post Comb, Low	Post Comb, Med	Post Comb, High	Pre Comb, Low	Pre Comb, Med	Pre Comb, High	
Main L	Expected $\alpha/\$$, using log average probabilities	Post Comb, Low	Post Comb, Med	Post Comb, High	Chem Loop, High	Pre Comb, High				
High Risk	$\alpha/\$$	Chem Loop, Low	Chem Loop, Med	Pre Comb, Low	Chem Loop, High	Post Comb, Low	Pre Comb, Med	Post Comb, Med	Pre Comb, High	Post Comb, High
Low Risk	Probability/\$, using average probabilities	Post Comb, Low	Chem Loop, Med	Chem Loop, High	Post Comb, Med	Post Comb, High	Pre Comb, Low	Pre Comb, Med	Pre Comb, High	
Low Risk L	Probability/\$, using log average probabilities	Post Comb, Low	Post Comb, Med	Post Comb, High	Pre Comb, High	Chem Loop, High				

Table 7: Alternative Funding Orders

We see that Chemical looping is funded first in two of the orders; whereas post-combustion comes first in the other three. Chemical looping has the highest impact, if successful, and so comes first in the high risk order. When using average probabilities, the probability of success is relatively high and the investment is low, so it has a high expected impact per dollar invested. However, when using log averages, the probability of success is low, therefore it does not rank high in that order. While post combustion has the smallest impact overall, it has very high probabilities of success, and so ranks high in a number of the lists.

Note that in some of the funding orders some of the projects are not listed at all (such as chemical looping low in the Main order). This is because they are not efficient: the appropriate metric is lower for a low funding

level than for a high funding level. For example, for chemical looping, the expected shift per dollar invested, $\frac{E[\alpha]}{I}$ is .001; .003. .002 for the low-, medium-, and high- funding trajectories. This would imply that we would first fund the medium, then the high, then the low. This makes no sense, and so the low funded project is simply not funded. Another way to think about this, is that there is increasing returns to scale between the low and medium project, therefore the low project is not efficient. On the other hand there are decreasing returns between the medium and low. Whether it is optimal to invest in the medium or low depends on the marginal value of the projects.

One salient point is that research into pre-combustion capture (that is combined with IGCC) is rarely higher on the lists than research into post-combustion capture (which is usually combined with pulverized coal). The reason is that, while the shift to the MAC is better *given success*, the probabilities of success that we elicited were uniformly low. We do note above, however, that there may be increasing returns to pre-combustion research.

In Figure 10 we show the expected return for the Main and High Risk funding orders (the Low Risk funding order is almost identical to the Main). The left panel uses the averaged probabilities; the right panel uses log-averaged probabilities. We have also factored in the probability of viability from the NAS study. We have assumed that the probability is 0.66 at zero funding; 0.77 at our maximum funding of \$960M; and linearly interpolated between the two points.

Note that there is a sharp bend for the funding orders in the left panel. When such an “elbow point” exists, it provides us with a portfolio that is optimal for a wide range of estimates of the value of pivoting the MAC down. The point where the bend is corresponds to funding Chemical looping at the high level and Post-combustion at the low level; and possibly Pre-combustion at the low level in the high risk case.

In the right panel, the Main-L funding order (and the Low Risk-L not shown) shows almost linear returns to R&D. If we assume that *no government funding* gives a probability of zero, then the elbow point would be at a low investment in Post-combustion. If, on the other hand, we assume that the curve is smooth, and that *no government funding* would give an expected alpha of about 0.08, then the elbow is at *no government funding*. The high risk curve is quite different. The sharp elbow here would indicate a high investment in Chemical looping and low investments in both pre- and post- combustion, the same as implied by the left panel.

While these curves incorporate the experts’ probabilities, they do not explicitly represent uncertainty in the

Figure 10: R&D Returns in terms of expected pivot to the MAC

Figure 11: Returns to R&D, by expert

efficacy of R&D. One way to include uncertainty in models is to separate out the individual experts. In this case, only two of the experts had probabilities over all three technologies. Figure 11 shows the expected pivot for Experts 1 and 3, using the Main funding order. We could put a $1/2$ probability on each curve, and use that to determine the optimal investment under uncertainty, or to determine the value of better information.

In Figure 12 we show how the probability of success is impacted by investment. We show the probability of success for each project as a function of funding for the main funding order, using averaged probabilities. We also show the probability of achieving success in at least one of the projects. We have incorporated the uncertainty over the long term viability of CCS into this figure. If the Main funding order is followed, then the first two investments give some probability of high success; the next two investments give some probability of low success;

Figure 12: Probability of success as a function of investment

and the next two give some probability of medium success. The last investment, a high investment into Post-combustion, gives a marginally higher probability of low success. This data can be used in models that represent R&D as increasing the probability of success (for example see [7][8]). This may be the most straightforward way to implement the randomness in a wide variety of models.

5 Discussion

We have performed expert elicitations on the potential for advancement in carbon capture technologies; implemented the definitions of technical endpoints into an Integrated Assessment Model in order to determine how the technologies are likely to impact the climate change problem, if they are successful; and presented results on the probabilistic relationship between R&D funding and climate change impacts.

The analysis has highlighted three key areas of concern, where the collection of more information may be of high value. First, there is disagreement over the probability of success of Post-combustion. If successful, this appears to lead to a significant impact on costs, pivoting the MAC down by 21%. Our experts, however, disagree over likelihoods, with probabilities of success at low funding levels ranging between 12% and 98%. Moreover, the NAS study put this level of success in their “failure” category. That is, they appear to assume that a carbon capture technology equivalent to our definition of success for Post-combustion will exist with very high probability without any government funding at all. Moreover, the baseline assumptions in the CCTP [14], for example, are

much more aggressive than our definition. Thus, a large part of the climate policy community is taking as given a technology which appears to be at least somewhat controversial. This suggests that it would benefit the carbon capture community to clearly discuss the hurdles for this technology. From the rationales, and from our other work, we believe that the significant disagreement is about reaching specified cost goals. On the other hand, our results indicate that achieving a particular cost goal may not be as important as simply having a viable technology.

Second, there is disagreement over the probability of success for Chemical looping. Because this technology is relatively cutting edge, and because carbon capture is very topical at the moment, we had difficulty finding experts to assess this. The differences between the two experts in their assessments are extreme, ranging from a probability of less than 1% to a probability of 84% at the high end of funding. This disagreement is particularly important as it drives key differences in the funding orders.

These two disagreements together lead to very different focal portfolios when we combine probabilities in different ways. When probabilities are averaged, or when a very high-risk, high-payoff strategy is preferred, then a portfolio of aggressive investments in chemical looping, combined with moderate investments in pre- and post- combustion appears fairly robust. On the other hand, when probabilities are calculated through log-odds averaging, the most robust policy appears to be *no government funding*.

Third, it is crucial that we understand the likelihood that, given technological success, CCS will be implemented widely. We have seen from our MACs that simply having CCS as an option to combat climate change is quite valuable. Yet, if the likelihood of implementing it is not high, then it reduces the attractiveness of a broad R&D investment in this technology (and increases the importance of pursuing other lines of research). The NAS study made a first attempt to assess this likelihood, but it is a very complicated question. It involves technical questions about the viability and long term security of geologic storage; plus a range of non-technical issues and social preferences. Thus, we recommend that a workshop or series of workshops be held to assess the probability of long term viability, along with a more in depth elicitation over the probability of success of the capture technologies.

Because this was an initial look at the range of current opinion, it is important for us to conclude with lessons learned and recommendations for future work along these lines. With relatively small numbers of experts, es-

pecially for chemical looping, results should be interpreted cautiously. The mean probabilities from such small samples are sensitive to the exact set of experts used; on the other hand, prior work has shown that the incremental value of adding another expert decreases significantly after 3-4 experts [47]. Thus, our results may be most profitably viewed as indicative of the level of agreement or disagreement about the prospects for various technologies. It is doubtful that the path to higher quality estimates is merely querying more experts (except in the case of chemical looping); rather, given the type of disagreements about both cost-targets and fundamental viability of some technologies, it would be to have a richer dialogue between the experts in order to get as close as possible to the underlying reasons for disagreement.

We used a mix of face-to-face conversations, phone calls and email to communicate with experts, and included a round of feedback to share viewpoints among experts. With greater resources for engagement of experts, it would be desirable to hold a multi-day workshop or series of workshops in which experts first meet to air issues, jointly clarify assumptions and definitions and structure event trees, then work individually with assessors to provide individualized probability assessments, and finally discuss where these assessments differ before finalizing the assessments. Because it was especially hard to agree on the prospects for meeting cost targets, preparation for the assessments might include pre-reading material on industry learning curves, participation of industry experts with particular experience on the issue of manufacturing cost, and perhaps structuring event trees to include the response of industry so that the circumstances leading to potential industrial scale cost reduction are less ambiguous. This should reduce the variance among assessments and improve the consistency of their rationales. With more time and resources available, we might also enrich the assessments in other ways. For example, rather than formulating a single definition of success, one could describe several different levels of success (or ranges) and assess probabilities of reaching each of them.

This paper is a first step toward treating CCS research activities as a portfolio in which risk and return should be balanced. We have found that experts indeed do have opinions about risk and return that can be incorporated in economic models. The economic implications of these opinions are illuminating, and also demonstrate the importance of additional dialogue to synthesize the knowledge of the research community with regard to society's management of this portfolio.

Acknowledgement 1 *This research was partially supported by the Office of Science (BER) U.S. Department of*

Energy, Grant No.DE-FG02-06ER64203. The authors gratefully acknowledge the contributions of participating technical experts Richard Doctor of Argonne National Laboratory; Barry Hooper of Cooperative Research Centre for Greenhouse Gas Technologies; Wei Liu of Pacific Northwest National Lab; and Gary Rochelle of The University of Texas at Austin. We also thank Leon Clarke of the Joint Global Change Research Institute for his work on the MiniCAM results.

References

- [1] Erin Baker and Kwame Adu-Bonnah. Investment in risky R&D programs in the face of climate uncertainty. *Energy Economics*, 30:465–486, 2008.
- [2] Erin Baker, Haewon Chon, and Jeffrey Keisler. NuclearPower: Combining expert elicitations with economic analysis to inform climate policy. *The Energy Journal*, Under Review, 2008.
- [3] Erin Baker, Haewon Chon, and Jeffrey Keisler. Advanced solar R&D: Combining economic analysis with expert elicitations to inform climate policy. *Energy Economics*, In Press, Forthcoming.
- [4] Erin Baker, Leon Clarke, Jeffrey Keisler, and Ekundayo Shittu. Uncertainty, technical change, and policy models. Technical Report 1028, College of Management, University of Massachusetts, Boston, 2007.
- [5] Erin Baker, Leon Clarke, and Ekundayo Shittu. Technical change and the marginal cost of abatement. *Energy Economics*, 30, 2008.
- [6] Erin Baker, Leon Clarke, and John Weyant. Optimal technology R&D in the face of climate uncertainty. *Climatic Change*, 78, 2006.
- [7] G. J. Blanford. *Technology Strategy for Climate Change: Optimal R&D Investment and its Interaction with Abatement Policy*. PhD thesis, Stanford University, 2006.
- [8] Christoph Bohringer and Thomas F. Rutherford. Innovation, uncertainty and instrument choice for climate policy. 2006. Working Paper. <http://rockefeller.dartmouth.edu/assets/pdf/Rutherford.pdf>.
- [9] Yann Bramoulle and Lars J. Olson. Allocation of pollution abatement under learning by doing. *Journal of Public Economics*, 89:1935–1960, 2005.

- [10] A.S. Brenkert, S. Smith, S. Kim, and H. Pitcher. Model documentation for the MiniCAM. Technical Report PNNL-14337, Pacific Northwest National Laboratory, 2003.
- [11] John F. Clarke and Jae A. Edmonds. Modeling energy technologies in a competitive market. *Energy Economics Journal*, 15:123–129, 1993.
- [12] Leon Clarke, James Edmonds, Henry Jacoby, Hugh Pitcher, John Reilly, and Richard Richels. Scenarios of greenhouse gas emissions and atmospheric concentrations. Final Report <http://www.climatechange.gov/Library/sap/sap2-1/finalreport/sap2-1a-final-all.pdf>, U.S. Climate Change Science Program, 2007.
- [13] Leon Clarke, Marshal Wise, Son Kim, Steve Smith, Joshua Lurz, James Edmonds, and Hugh Pitcher. Model documentation for the minicam climate change science program stabilization scenarios: Ccsp product 2.1a. Technical report, Pacific Northwest National Laboratory, July 2007.
- [14] Leon Clarke, Marshall Wise, Marylynn Placet, R Cesar Izaurrealde, Joshua Lurz, Son Kim, Steve Smith, and Allison Thomson. Climate change mitigation: An analysis of advanced technology scenarios. Technical Report PNNL-16078, Pacific Northwest National Laboratory, 2006.
- [15] R.T. Clemen and R.L. Winkler. Combining probability distributions from experts in risk analysis. *Risk Analysis*, 19:187–203, 1999.
- [16] J. David and H. Herzog. The cost of carbon capture. Presented at the Fifth International Conference on Greenhouse Gas Control Technologies, Cairns, Australia, August 13 - August 16, 2000.
- [17] DOE. *Carbon Sequestration Technology Roadmap and Program Plan*. United States Department of Energy, Washington D.C., 2007.
- [18] J.J. Dooley, R.T. Dahowski, C.L. Davidson, M.A. Wise, N. Gupta, and S.H. Kim. *Carbon Dioxide Capture and Geologic Storage: A Core Element of a Global Energy Technology Strategy to Address Climate Change*. Global Energy Technology Strategy Program, 2006.
- [19] Paul B. Downing and Lawrence J. White. Innovation in pollution control. *Journal of Environmental Economics and Management*, 13:18–29, 1986.

- [20] J.A. Edmonds, J.F. Clarke, J.J. Dooley, S.H. Kim, and S.J. Smith. Stabilization of CO₂ in a B2 world: Insights on the roles of carbon capture and storage, hydrogen, and transportation technologies. In John Weyant and Richard Tol, editors, *Special Issue, Energy Economics*. 2005.
- [21] Carolyn Fischer, Ian W.H. Parry, and William A. Pizer. Instrument choice for environmental protection when technological innovation is endogenous. *Journal of Environmental Economics and Management*, 45:523–545, 2003.
- [22] L. Goulder and S. Schneider. Induced technological change and the attractiveness of CO₂ abatement policies. *Resource and Energy Economics*, 21:211–253, 1999.
- [23] Lawrence Goulder and Koshy Mathai. Optimal CO₂ abatement in the presence of induced technological change. *Journal of Environmental Economics and Management*, 39:1–38, 2000.
- [24] Harold Hotelling. The economics of exhaustible resources. *Journal of Political Economy*, 39:137–175, 1931.
- [25] R. A. Howard and J. E. Matheson. When is it appropriate to combine expert judgments? *Decision Analysis*, 2, 2005.
- [26] G. S. Esber III. *Carbon dioxide capture technology for the coal-powered electricity industry : a systematic prioritization of research needs*. MIT, 2006.
- [27] IPCC. *Special Report on Carbon Dioxide Capture and Storage*. Cambridge University Press, New York, NY, 2005.
- [28] Melanie D. Jensen, Mark A. Musich, John D. Rubya, Edward N. Steadman, and John A. Harju. Carbon seperation and capture. Technical report, Energy and Environmental Research Center, National Energy Technology Laboratory, June 2005.
- [29] Chulho Jung, Kerry Krutilla, and Roy Boyd. Incentives for advanced pollution abatement technology at the industry level: An evaluation of policy alternatives. *Journal of Environmental Economics and Management*, 30:95–111, 1996.
- [30] J. Katzer, E. J. Moniz, J. Deutch, S. Ansolabehere, J. Beer, and et Al. The future of coal: An interdisciplinary MIT study. Technical report, Massachusetts Institute of Technology, 2007.

- [31] D. W. Keith. When is it appropriate to combine expert judgments? *Climatic Change*, 33, 1996.
- [32] B. Metz, O. Davidson, H.C. de Coninck, M. Loos, and L.A. Meyer, editors. *IPCC Special Report on Carbon Dioxide Capture and Storage*, Cambridge U.K., 2005. Cambridge University Press.
- [33] Scott R. Milliman and Raymond Prince. Firm incentives to promote technological change in pollution control. *Journal of Environmental Economics and Management*, 17:247–265, 1989.
- [34] Juan-Pablo Montero. Permits, standard, and technology innovation. *Journal of Environmental Economics and Management*, 44:23–44, 2002.
- [35] M. G. Morgan and D. W. Keith. Subjective judgments by climate experts. *Environmental Science and Technology*, 29, 1995.
- [36] National Research Council. *Prospective Evaluation of Applied Energy Research and Development at DOE (Phase Two)*. The National Academies Press, Washington D.C., 2007.
<http://www.nap.edu/catalog/11806.html>.
- [37] Gregory F. Nemet and Erin Baker. Demand subsidies versus R&D: Comparing the uncertain impacts of policy on a pre-commercial low-carbon energy technology. *The Energy Journal*, Under Review.
- [38] Ian Parry. Pollution regulation and the efficiency gains from technological innovation. *Journal of Regulatory Economics*, 14:229–254, 1998.
- [39] S. C. Peck and Y. H. Wan. Analytic solutions of simple greenhouse gas emission models. In E.C. Van Ierland and K. Gorka, editors, *Economics of Atmospheric Pollution*, page Chapter 6. Springer Verlag, Berlin, 1996.
- [40] K. Riahi, E.S. Rubin, M.R. Taylor, L. Schrattenholzer, and D. Hounshell. Technological learning for carbon capture and sequestration technologies. *Energy Economics*, 2004.
- [41] Knut Einar Rosendahl. Cost-effective environmental policy: Implications of induced technological change. *Journal of Environmental Economics and Management*, 48:1099–1121, 2004.
- [42] Edward S. Rubin, Margaret R. Taylor, Sonia Yeh, and David A. Hounshell. Learning curves for environmental technology and their importance for climate policy analysis. *Energy*, 29:1551–1559, 2004.

- [43] Ahti A. Salo and Derek W. Bunn. Decomposition in the assessment of judgmental probability forecasts. *Technological Forecasting and Social Change*, 49:13–25, 1995.
- [44] Carl S. Spetzler and Carl-Axel S. Stael von Holstein. Probability encoding in decision analysis. *Management Science*, 22:340–358, 1975.
- [45] Detlof von Winterfeldt and Ward Edwards. *Decision Analysis and Behavioral Research*. Cambridge University Press, Cambridge, U.K., 1986.
- [46] John P. Weyant and Jennifer Hill. Introduction and overview. *The Costs of the Kyoto Protocol: A Multi-Model Evaluation. Special Issue of The Energy Journal.*, pages vii–xliv, 1999.
- [47] Robert L. Winkler and Robert T. Clemen. Multiple experts vs. multiple methods: Combining correlation assessments. *Decision Analysis*, 1:167–176, 2004.