

GHGT-11

Rethinking CCS – Developing quantitative tools for designing robust policy in face of uncertainty

Jan Eide^a, Howard Herzog^{a,*}, Mort Webster^b

^aMIT Energy Initiative, 77 Massachusetts Avenue, Cambridge MA 02139, USA

^bMIT Engineering Systems Division, 77 Massachusetts Avenue, Cambridge MA 02139, USA

Abstract

Lack of climate policy and CO₂ markets along with a global economic slowdown suggest that we need to rethink our approach to demonstrating CCS at a commercial scale. Austerity measures make it likely that public funding will be tight in coming years, and there is a striking need to ensure that limited funds are spent optimally. Quantitative tools exist for aiding decision making under uncertainty, yet few of them have been applied to build a model that can help answer the question of what is the optimal allocation of a given amount of money across a portfolio of demonstration projects that maximizes learning about CCS. Developing such a model is the goal of this paper and we employ the model to assess the proper role of Enhanced Oil Recovery (EOR) in a CCS demonstration portfolio. We find that if we want to maximize learning, a CCUS-only (CCS + EOR) approach to developing CCS as a mitigation technology would only be advisable if there was little uncertainty in non-EOR storage. As we believe that this condition is unlikely to be true, we suggest that U.S. policy makers should be particularly cautious in relying on a CCUS-only approach to CCS development. Nonetheless, we also find that a portfolio consisting of a mix of CCS and CCUS projects can be an effective strategy in a number of situations, notably if EOR can teach us important lessons about non-EOR storage.

© 2013 The Authors. Published by Elsevier Ltd.
Selection and/or peer-review under responsibility of GHGT

Keywords: CCS policy; CCUS; R&D portfolio analysis; Dynamic programming; Bayesian learning

1. Introduction

Rising CO₂ concentrations in the atmosphere could potentially have dramatic consequences for the world's climate in the coming decades (IPCC, 2007) [1]. With its promise of nearly CO₂-free electricity from fossil fuel sources, carbon capture and storage (CCS) has been viewed as an important option to

* Corresponding author. Tel.: +1-617-253-0688; fax: +1-617-253-8013.
E-mail address: hjherzog@mit.edu.

consider for reducing future emissions. Consequently, over the past decade many governments laid out aggressive roadmaps for CCS development and deployment. Yet no global emissions reduction agreement has been signed, and the failure of passing a cap and trade bill through the U.S. Congress has made it highly unlikely that the primary driver for CCS projects, climate policy, will be in place in the U.S. this decade. Consequently there will be no climate markets for CO₂, and coupled with a global economic slowdown this new reality suggests that we need to rethink our approach to demonstrating CCS at a commercial scale.

Because the long-run cost and performance of CCS is uncertain, there is significant value in exploring the potential of this mitigation technology, even despite the worsened outlook for climate policy. The next several decades provide a valuable opportunity for learning. If public policy in the future calls for large-scale emission reductions, it is important to advance the state of readiness for CCS technologies and to determine what role they can play in a climate mitigation portfolio.

Austerity measures in many developed countries will likely result in limited funds for large-scale demonstration projects, making it even more important to ensure that resources are strategically allocated to achieve the highest return. As governments generally invest in R&D and demonstration projects to gain knowledge, we make the assumption that policy-makers want to invest in a way that maximizes knowledge acquisition. Knowledge is of course a vague term, and can include better physical understanding of reservoirs (e.g., regarding reservoir leakage), ways to reduce costs, developing new technologies, etc. All this helps us reduce uncertainty as we determine the role of CCS in a future low-carbon energy portfolio. As “knowledge” is hard to model quantitatively, we will assume that reducing uncertainty correlates to acquiring knowledge. As a first step we will simplify the knowledge acquisition problem by just modeling the uncertainty in cost, with cost being a proxy for a wide range of technical and economic issues. Policymakers will need to determine the optimal allocation of a given amount of money across a portfolio of demonstration projects that minimizes the uncertainty in the cost of CCS.

Arguably this is a very complex problem, not least because we have no accurate way of determining the effect of any demonstration project on the acquisition of knowledge or the reduction of uncertainty. Using a number of simplifying assumptions we can nonetheless develop a quantitative optimization model of project selection under uncertainty. Specifically, we apply a dynamic portfolio optimization framework with Bayesian learning to assess how different carbon capture projects reduce uncertainty about CCS as a mitigation technology. The focus of this paper is not to contribute to the state of the art of mathematical portfolio analysis, but rather applies the methods in an illustrative example to provide new insight about the future path of CCS policy in the U.S. and globally.

The model will be used to answer a specific and timely question: What is the appropriate role for CCUS in a CCS portfolio? With a large potential market for captured CO₂, Enhanced Oil Recovery (EOR) has gained prominence in recent years with U.S. policy makers due to the lower cost of CCUS projects. Yet CCUS is not an end-goal in itself, and the assumption has been that experience with CCUS will spill over to CCS. We will use our optimization model to gain insight about whether CCUS is conducive to gaining knowledge about CCS, and if so, under what conditions?

The structure of the paper is as follows. In Section 2, we review the literature in two related areas, climate and energy economics, and operations research. In Section 3, we frame the decision problem studied, and describe the modeling methodology. The modeling results are presented in Section 4. Section 5 concludes with a discussion of the insights for energy R&D policy.

2. Background

This paper builds on two distinct areas of study in assessing the optimal path for CCS development: energy and climate economics, and operations research and dynamic portfolio optimization. Economists have conducted numerous studies on the question of how to optimally spend public R&D funds to reduce the cost of future emission reductions; see Gillingham et al. (2008) [2] for a good overview of different approaches used to model technological change. One approach is to incorporate technological change in integrated general-equilibrium models, such as WITCH (Bosetti et al., 2011 [3]) and DICE (Nordhaus, 2002 [4]). The integrated modeling approach captures the dynamics between R&D investments, the evolution of technologies, and how such changes affect the relative economics of one mitigation technology over another in meeting an emission reductions target. As pointed out by McJeon, et al., (2011) [5], a key challenge with many integrated assessment models is that they generally use representative scenarios for how one technology evolves, and thus do not consider the broad range of possible paths that a technology can take in response to a demonstration or R&D program. Webster et al. (2008) [6] explore this issue by applying Monte Carlo simulation, and McJeon, et al., (2011) [5] uses a combinatorial approach on the U.S. government's Climate Change Technology Program (CCTP)'s assumptions of technical development. However, Monte Carlo approaches do not address decision under uncertainty, which is our focus here.

The current literature on incorporating technological learning into a modeling framework generally distinguishes between learning-by-doing and learning-by-searching (R&D-based approaches). Learning-by-searching models of the R&D process typically rely on the concept of “knowledge capital” in their representation of endogenous learning (e.g., Popp, 2004 [6]; Bosetti et al., 2011 [3]). The concept of learning-by-doing, or learning rates (e.g., McDonald & Schrattenholzer, 2001 [8]; Rubin et al., 2004 [9]) focuses on reductions in technology cost that occur as a function of cumulative investment or cumulative production during the commercial phase of technology development.

Recent R&D portfolio analysis, e.g., Ringuest et al. (2004) [10] has built upon the rich tradition of financial portfolio optimization going back to Markowitz (1959) [11] and Samuelson (1969) [12]. The classic decision criterion is the mean-variance analysis, by which the optimal portfolio either minimizes variance for a given mean return, or maximizes the mean return for a given variance (Chien, 2002) [13]. The literature on Bayesian portfolio analysis (e.g., Avramov & Zhou [14], 2010; Zellner & Chetty, 1965 [15]) explicitly models how observations over time lead to revised beliefs about uncertainty.

However, optimizing a project portfolio is fundamentally different from a financial portfolio optimization. As pointed out by Vilkkumaa et al. (2012) [16], decisions are not continuous (i.e., you cannot build a fraction of a CCS project), and markets do not determine their price or value. Some studies have adapted the traditional portfolio optimization framework to consider energy project R&D decisions. Chao et al. (1990) [17] use a two-stage model with discrete investment amounts to explore how to allocate initial exploratory R&D across several candidate energy technologies. Guo (2012) [18] extends the portfolio selection framework to explicitly include the value of endogenous learning (i.e., learning that only occurs with investments and scales with the amount invested), and demonstrates the model with a two-technology decision problem over many stages.

We adapt a multi-stage R&D investment portfolio optimization model to explore the optimal allocation across four types of CCS projects which result in differing amounts of learning (reduction in the uncertainty) about capture and sequestration. Our model, described in the next section, is closest to that of Guo (2012) [18], but is applied to a specific question within CCS technology development.

3. Methodology

To explore the value of CCUS as a strategy for CCS technology development, we construct a stylized decision problem. This simple example will enable us to highlight the trade-offs across a wide range of possible situations. The biggest benefit of this exercise is not to generate a detailed allocation of R&D funds, but to help one better understand how to think about this problem and gain insights in making the allocation decisions.

Consider a decision maker with a fixed budget every period that can be allocated over a number of CCS demonstration projects. His goal is to reduce the uncertainty in project costs (our proxy for gaining knowledge). Some projects cost more, but result in more useful information for reducing the uncertainty in CCS costs. Also, the cost of each individual project is variable relative to the average cost, which reduces the information learned from a single project. After observing the cost of the chosen projects, the decision-maker again chooses new projects to invest in for the next period, and the process repeats. The decision problem in any period is how to allocate funding across project types in order to maximize the reduction in cost uncertainty by the final period.

We want to predict the cost of one project ahead of time and our prediction relies on our understanding of the average cost of all projects. Yet the resulting cost estimate for a single project will be uncertain for two reasons. First there is uncertainty regarding the average cost due to a lack of knowledge and experience. In the model, this is referred to as “uncertainty”, and it will decrease as we gain knowledge from demonstration plants. Second, although we can use the average cost to predict the cost of individual projects by considering site-specific factors, all the heterogeneities of individual projects cannot be accounted for ahead of time. Individual project costs will therefore vary around the average, even if we have tried to account for project-specific heterogeneities. In the model this is referred to as “variability”, and it will persist even when we have enough knowledge to determine the average cost with confidence. For example, when building a refinery, you can factor site-specifics such as tax rate, land cost, labor cost etc. into the cost ahead of time. Yet even if there are decades of experience and data on the average cost, individual project costs can still be a bit lower or a bit higher than anticipated due to the variability that can be associated with the heterogeneities of specific projects.

Below, we formalize the model of this decision problem, and describe the assumptions made for this illustration.

3.1 Dynamic Programming Formulation

We frame the decision problem and solve it using stochastic dynamic programming. Dynamic programming (DP) provides a structure for solving multi-stage sequential decision problems under uncertainty. Rather than solve the entire problem at once, which is generally prohibitively large, we decompose the problem and solve iteratively for the optimality conditions at every decision stage.

To formalize the problem, we use a Bayesian approach to model the uncertainty and learning. We assume that the cost of any CCS project is distributed normally as

$$C \sim N(\mu, \sigma)$$

where μ is the mean or average cost and σ is the standard deviation, which represents the variability of projects. Because the mean cost of CCS projects are uncertain, we represent that uncertainty as a probability distribution for the average cost μ , and assume that it is distributed as

$$\mu \sim N(m, s)$$

where m is the mean, or current “best-guess” for the average cost, and s is the standard deviation, which represents our current uncertainty in the average cost of CCS projects. After each new CCS investment, the actual cost for that project will be observed, and updating m and s according to Bayes rule reduces the uncertainty in future CCS project costs.

To capture the key features of the current debate over whether to invest in demonstration projects that capture carbon from high-purity sources and/or use the carbon for EOR, we further disaggregate the costs of each CCS project into the sum of two components: the cost of capture C^c and the cost of sequestration C^s .

$$C = C^c + C^s$$

The uncertainties in the capture cost and the sequestration cost are represented separately, and each observed cost from an investment updates both uncertainties:

$$\begin{aligned} C^c &\sim N(\mu^c, \sigma^c), \mu^c \sim N(m^c, s^c) \\ C^s &\sim N(\mu^s, \sigma^s), \mu^s \sim N(m^s, s^s) \end{aligned}$$

We simplify the range of possible CCS demonstration projects into four possible types (Figure 1). The carbon capture can occur within a high-purity industrial process or within an electricity generation facility. The carbon can then be used for EOR, or sequestered in a non-EOR reservoir such as a saline aquifer. The four possible project types are high-purity capture and EOR (HP-CCUS), high-purity capture with non-EOR sequestration (HP-CCS), power plant capture with EOR (CCUS), or power plant capture with non-EOR sequestration (CCS). However, if the ultimate goal is to use CCS for climate mitigation, then power generation capture and non-EOR storage will need to play major roles. We are therefore interested in reducing the uncertainty in this capture and storage method. We assume that the cost of capture from a high-purity source is less than capture in a power plant.

$$C^{HP} < C^c,$$

and the cost of storage for EOR is less than non-EOR storage

$$C^{EOR} < C^s.$$

	High-purity capture	Power generation capture
EOR storage	<p>HP-CCS</p> <p>Learning:</p> <ul style="list-style-type: none"> - Less learning about capture - Less learning about storage <p>Cost</p> <ul style="list-style-type: none"> - Least expensive 	<p>CCUS</p> <p>Learning:</p> <ul style="list-style-type: none"> - Learn about capture - Less learning about storage <p>Cost</p> <ul style="list-style-type: none"> - Expensive
Non-EOR storage	<p>HP-CCS</p> <p>Learning:</p> <ul style="list-style-type: none"> - Less learning about capture - Learn about storage <p>Cost</p> <ul style="list-style-type: none"> - Less expensive 	<p>CCS</p> <p>Learning:</p> <ul style="list-style-type: none"> - Learn about capture - Learn about storage <p>Cost</p> <ul style="list-style-type: none"> - Very expensive

Fig. 1: Carbon capture project types

We now formalize the dynamic programming decision problem. The decision-maker's objective is to minimize the uncertainty in CCS costs by the terminal period T :

$$\min s_T$$

To simplify the example here, we assume that the decision-maker can only choose up to two projects in each period t :

$$a_{t,1}, a_{t,2} \in \{\text{HP-CCUS}, \text{HP-CCS}, \text{CCUS}, \text{CCS}\}$$

The state variable, which fully capture all relevant information about the evolution of the system up to period t , are the parameters that describe the cost uncertainties based on all projects observed up to this point. For this problem, the state variable x_t at period t is the vector:

$$x_t = \{m_t^c, s_t^c, m_t^s, s_t^s\}$$

The state transition equations describe how the state evolves as a function of the action chosen and the random variation that occurs. When the next set of two projects is chosen, the capture and sequestration costs for each project are drawn randomly from the current probability distribution. The observed costs of capture and sequestration, \widehat{C}^c and \widehat{C}^s , from each project are then used to update the parameters according to:

$$m_{t+1} = \frac{\sigma^2 m_t + s_t^2 \widehat{C}}{\sigma^2 + s_t^2}$$

$$s_{t+1}^2 = \frac{\sigma^2 s_t^2}{\sigma^2 + s_t^2}$$

To solve this problem using dynamic programming, we use backward induction to recursively solve the Bellman value function:

$$V_t = \min_{a_{t,1}, a_{t,2}} E\{S_T^2\}(a_{t,1}, a_{t,2}, x_t)$$

For the example shown below, we assume a finite horizon problem with three periods,
 $t = \{1, 2, 3\}$,
 each representing 10 years.

Finally, we add to the model the assumption that the observed cost of capture from a high-purity project is less useful for reducing the uncertainty in capture costs from power plants, and that the observed cost of storage from an EOR project is less useful in reducing the uncertainty in non-EOR storage costs. We model this reduced learning by using a 2x1 weighting vector $w \in [0, 1]^2$. w describes how much of the learning obtained should actually be considered in the updated posterior distributions. For example w^s will be 0 if we assume no transferable learning from EOR to non-EOR storage. If we assume half the learning to be transferable w^s will be 0.5.

The reward function for this example depends solely on minimizing the uncertainty in the final period T , and does not consider any time-value of learning. This assumption is for simplicity, and will be relaxed in future work.

3.2 Cost Assumptions

Cost assumptions must be treated with caution. Flyvbjerg et al. (2003) [19] showed that costs are generally underestimated for large and complex infrastructure projects, and there is reason to be similarly cautious for CCS projects. The inputs used should therefore be regarded as more of a representation of a stylized type of project, highlighting differences between the four quadrants in Figure 1, rather than an accurate prediction of what future costs will be. In our model we report capture costs in \$/tonne avoided whereas storage costs are reported in \$/tonne captured[†].

The cost of capture from power plants has been referenced thoroughly in the literature. IEA (2011) [20] examined cost studies from eight different organizations with avoided costs ranging from \$40-\$69/tonne CO₂. Recent data in ZEP (2011) [21] also suggests an avoided cost of around \$44/tonne CO₂. Given the inherent uncertainty of cost estimates we nonetheless would like to model a greater range of uncertain costs, particularly for demonstration projects. Estimates of first of a kind avoided costs at Norway's Mongstad project was estimated by one report at between approximately \$228-\$395/tonne CO₂ (Klif, 2010 [22]). Although these costs are probably not representative of likely future average Nth of a kind costs, they do highlight the significant uncertainty that surrounds the cost of demonstration projects. The Gaussian probability distribution that we believe captures this range of uncertainty is one with mean values ranging from \$40-\$160/tonne CO₂ avoided, and a maximum standard deviation of \$15.5/tonne (equivalent to +/- \$30/tonne).

The symmetric Gaussian distribution may not necessarily be the most appropriate choice for modeling uncertain costs for large engineering projects. Actual costs are probably more likely to be higher than anticipated, rather than lower. That will not be the case for a Gaussian distribution with a certain mean

[†] Converting from tonnes of CO₂ avoided to tonnes of CO₂ captured is done by dividing by $(1 - \frac{\text{energy penalty}}{\text{capture percentage}})$. For an energy penalty of 25% and 90% capture this is equivalent to dividing by a factor of 0.72. Therefore, 100 tonnes avoided is equivalent to 139 tonnes captured. Also \$100/tonne captured equals \$72/tonne avoided.

cost: the probability that actual costs will be lower than the mean is equal to the probability that cost will be higher than the mean. A more appropriate model of cost uncertainty would be some positively skewed distribution with a right-hand “tail” for very high average costs. As more projects are realized, the “tail” of the distribution will shrink, and the distribution will converge on the actual mean cost. Future versions of the model will be expanded to include for different distribution families.

Data from Alstom (2011) [23] suggests a non-EOR storage cost of around \$10/tonne CO₂ captured for onshore storage, with an additional \$5/tonne CO₂ captured for transport. Yet due to lack of experience with large-scale injection and long-term storage, it is more appropriate to use a range of storage costs. More specifically, an additional \$40/tonne CO₂ is added to account for any contingencies related to long-term monitoring and potential mediation of leaks. The Gaussian probability distribution that we believe captures this range of uncertainty is one with mean values ranging from \$10-\$55/tonne CO₂ captured, and a maximum standard deviation of \$7.5/tonne CO₂ (equivalent to +/- \$15/tonne CO₂). As with capture cost, some positively skewed distribution is likely a better choice to represent cost uncertainty and this will be addressed in future versions of the model.

We assume the cost of high-purity capture to be \$16/tonne CO₂ captured (equivalent to \$22/tonne CO₂ avoided), similar to OPEX costs at Sleipner, reported in IEA (2008) [24]. EOR storage cost is set at - \$15/tonne CO₂ captured[‡].

Standard deviations do not necessarily give an immediate, intuitive sense of the degree of uncertainty. The distance from the mean that samples can fall is a more intuitive metric. Table 1 therefore shows the 95% confidence interval for different standard deviations, rounded to the nearest integer. For example, if the standard deviation of the cost distribution is \$7.5/ton and the mean value \$25/ton, then, with 95% certainty, actual costs will fall in the interval [10\$/ton, \$40/ton].

Table 1 Relationship between standard deviation and uncertainty range

Standard deviation (\$/tonne)	Deviation from mean value (\$/tonne)
7.5	+/- 15
15.5	+/- 30

Demonstration projects are likely to cost more than future expected Nth plant costs. In the results section we will therefore use a base case of an expected mean cost of capture of \$100/tonne CO₂ avoided and an expected mean storage cost of \$25/tonne CO₂ captured. These costs are significantly higher than the numbers in (IEA, 2011) [20], but we feel that a conservative estimate, rather than a too optimistic one, should be the basis of CCS policy.

To calculate the total lifetime cost of a CCS project we consider a 500 MW net supercritical coal plant with an emissions rate of 830 g CO₂/kWh without CCS. Assuming a 25% energy penalty this results in an emission rate of 1107 g CO₂/kWh for the CCS plant. With 90% capture and a 75% capacity factor a total of 3.27 Mt CO₂ will be captured annually, although the amount of CO₂ avoided will only be 2.36 Mt.

[‡] Assuming 2% of oil price at \$70/bbl in Mcf, 18 ton/Mcf, subtracting \$10/ton for transport.

Consequently, total annual cost of a CCS project is the cost of capture and storage, in \$/ton avoided, multiplied by 2.36 Mt. The thirty-year net present value of costs is calculated using a 7% discount rate.

The degree of variability in capture and storage costs is hard to determine *ex-ante*, as we have not yet observed how far from the average the cost of individual projects will be. We believe that there will be significant variability in the cost of capture from project to project, but a lot of this can be accounted for, such as differences due to land cost, coal type, labor cost etc. Consequently we believe that the variability that we *cannot account for* will be small. However, for storage costs, it is likely to be harder to account for project heterogeneities. Site-specific particularities of geologic formations could impact costs in unpredictable ways, and consequently there will be a lot of variability in storage cost that we cannot account for. To illustrate the difference in capture and storage variability we choose a capture standard deviation of \$4/tonne CO₂ avoided (equivalent to +/- 8/tonne CO₂) and a storage cost standard deviation of \$8/tonne CO₂ captured (equivalent to +/- 16/tonne CO₂). It is important to acknowledge the lack of empiric data to test our assumptions. Nonetheless, as our modeling approach is meant for illustrative purposes, we use the above variability to illustrate our belief that the variability we cannot account for will be greater for storage cost than for capture cost.

Although we will analyze the effect of different budget levels per period in the result section we assume a baseline budget of \$8 billion per investment period. This amount counts both public and private funds made available for CCS demonstration. With a total of around \$180 million spent on capture and storage projects by the U.S. Department of Energy in 2011, and another \$165 million requested for 2013 (DOE, 2012 [25]), it seems that our estimate of public funds available per decade is reasonable.

3.3 The relative roles of uncertainty and variability

In our notation, the variability of any individual project relative to the mean cost of projects is σ , while the uncertainty in the average costs is s . As described above (see also Appendix 1), after the realized cost of a project is observed, the revised uncertainty s_{t+1} in average costs, expressed as a standard deviation, is given by:

$$s_{t+1} = \sqrt{\frac{\sigma^2 s_t^2}{\sigma^2 + s_t^2}} \quad (1)$$

Reorganizing (1) yields

$$\frac{s_{t+1}}{s_t} = \frac{1}{\sqrt{1 + \left(\frac{s_t}{\sigma}\right)^2}} \quad (2)$$

Because $\left(\frac{s}{\sigma}\right)^2$ will always be greater than zero, (2) is a monotonically decreasing function. Furthermore, (2) shows that s_{t+1} is always less than s_t , which is intuitive given that it represents our updated knowledge. When the ratio $\frac{s}{\sigma}$ is small, representing that our uncertainty in average costs is small relative to a large variability in any individual project, the learning effect will be small. Similarly, when this ratio is large, representing that our uncertainty in average costs is large while the individual variability is small, each observation will yield significant uncertainty reduction. If the distribution from which samples are drawn exhibits little variability (i.e., σ is small) then any sample drawn will likely be very close to the actual mean of the distribution and learning will be significant. On the other hand, if samples are drawn from a

distribution with significant variability, more samples will be needed to obtain a reasonable estimate of the mean. Each individual observation provides less information, and more samples will be needed in order to reduce the uncertainty. As illustrated graphically in Fig. 2, learning decreases with increasing variability (σ). A subtlety to be understood is that this treatment assumes that the entity observing the costs will know beforehand whether the variability is large or small. The representativeness of any observation of cost for the average value depends on the variability.




Fig. 2: Larger σ reduces learning

4. Results

The total demonstration budget plays a key role in determining the optimal portfolio of projects. If the budget is too low, only high-purity projects will be undertaken due not enough funds being available for power sector projects. If the budget is unlimited, only CCS projects will be undertaken due to their maximization of learning. The interesting dynamic nonetheless occurs from the situation in between these two extremes, where not enough money exists to only do CCS projects, but not so little as to rule them out completely. Also, for the results shown here, we assume that the variability in sequestration costs is greater than the variability in capture costs. This assumption is based on the hypothesis that the developing accurate cost estimate models for geologic storage is harder than it is to develop accurate cost estimate models for power plants.

For a first analysis, we assume that EOR storage does not reduce storage uncertainty and that high-purity capture does not reduce capture uncertainty (i.e. CCUS projects only reduce capture uncertainty, and HP-CCS projects only reduce storage uncertainty). We assume an expected mean capture cost (m_{capture}) of \$100/tonne CO_2 avoided and an expected mean storage cost (m_{storage}) of \$25/tonne CO_2 captured. Figure 3 shows the optimal investment strategy as a function of the absolute uncertainty (s) in capture costs and sequestration costs. To get an intuitive understanding of the relationship between standard deviations (i.e. uncertainty) and ranges of possible values refer to Table 1. Current proposals to shift investments to solely EOR storage (CCUS) are only optimal when sequestration costs are nearly certain (dark blue region along bottom). For relatively low uncertainty in sequestration costs, the optimal strategy is a mix

of one CCS and one CCUS (light blue region). For increasing standard deviations in storage costs, the preferred portfolio is one CCS project and one HP-CCS project (red region).




Fig. 3: Optimal project portfolios for large storage variability ($\sigma=\$8/\text{ton}$) and low capture variability ($\sigma=\$4/\text{ton}$)

Next we assume that EOR storage reduces storage uncertainty and that high-purity capture will reduce capture uncertainty (i.e. CCUS projects will reduce both storage and capture uncertainty, and HP-CCS projects will also reduce both capture and storage uncertainty). We model this with the weighted approach described in section 3. If EOR projects reduce the uncertainty in sequestration costs (with a weight of 0.8), the strategy (CCS, CCUS) is optimal for close to all ranges of storage uncertainty (Figure 4, left panel). If high-purity projects reduce the uncertainty in capture costs, a (CCS, CCUS) strategy is optimal over a smaller range of storage cost uncertainty (Figure 4, right panel). In general, the degree to which this boundary between strategies moves depends on how close the learning weight for EOR (HP) is to 1, and how close the learning weight for HP (EOR) is to zero. A key result is that even if the learning weight for EOR storage is 0.8, a portfolio of only CCUS projects is advisable only if storage uncertainty is effectively zero. Nonetheless, a high learning weight for EOR storage expands the light blue area so that a partial shift to CCUS is advisable for a greater number of situations.




Fig. 4 (a) Effect of learning from EOR ($w=[0.8,0]$); (b) Effect of high-purity learning ($w=[0, 0.8]$)

It is plausible that some experience with EOR storage would reduce uncertainty in non-EOR storage. However, it is harder to imagine that experience with high-purity capture projects would significantly reduce the uncertainty in power plant capture. We base this argument on the simple fact that high-purity capture of CO₂ has been done commercially for decades without having a discernible positive impact on reducing uncertainty about power plant capture costs.

5. Conclusion

The key insight of the model is that the relative amounts of uncertainty and variability will have a significant impact on determining optimal CCS portfolios, if the near-term objective is to gain knowledge and reduce uncertainty. The key insight is that the more variability there is, the harder it is to reduce uncertainty in the average cost. This is important for CCS, since variability will have a large impact on the difficulty of learning about the cost of storage and capture. We argue that accounting for project-specific heterogeneities in geologic storage is harder than accounting for project-specific heterogeneities in capture plants. The variability we cannot account for will therefore be greater for storage cost than for capture cost. Consequently we will need more observations to reduce average storage cost uncertainty than we will need to reduce average capture cost uncertainty. Since cost uncertainty is a proxy for lack of knowledge, we will need more storage demonstrations than capture demonstrations to fully develop CCS as a viable mitigation option.

With public funding likely to be limited in coming years, the simple, stylized example presented here provides valuable insight about the optimal allocation of funds across different demonstration projects. Notably, a CCUS-only approach (investing exclusively in power projects with EOR storage) to developing CCS as a mitigation technology would only be advisable if there was little uncertainty regarding non-EOR storage. Given our lack of experience with large-scale injection of CO₂ in geologic formations, this condition is unlikely to be true. U.S. policy makers should therefore be particularly cautious in suggesting a CCUS-only approach to CCS development.

However, a portfolio consisting of a mix of CCS and CCUS projects is an effective strategy to gain knowledge if capture uncertainty is high while non-EOR storage uncertainty is low and/or EOR storage projects can teach us important lessons about non-EOR storage.

Due to the larger absolute values of capture costs, it seems fair to assume that absolute capture uncertainty will be higher than absolute storage uncertainty. Yet it is unclear whether storage cost uncertainty is low enough to justify CCUS in the case where no learning from EOR storage is transferable to non-EOR storage. A partial move towards CCUS can however be advisable if we simultaneously also believe that EOR storage will provide significant knowledge for non-EOR storage.

It is clear that the important roles of variability, uncertainty, and learning from EOR storage in determining optimal demonstration portfolios have not been studied thoroughly by the CCS community. The results of this analysis suggest that more is warranted. Future extensions to this work include extending the model to consider combinations of more than two projects per period, and more extensive sensitivity analysis. Coupling the expanded model with better estimates of actual variability and uncertainty will provide important insight about the future path of CCS development.

Acknowledgements

This paper is based on research conducted at MIT from September 2011–October 2012. We would like to thank the MIT Carbon Sequestration Initiative for their financial support of this research effort.

References

- [1] IPCC. *Synthesis Report*. Intergovernmental Panel on Climate Change 2007.
- [2] Gillingham K, Newell R, Pizer W. Modeling endogenous technological change for climate policy analysis. *Energy Economics* 2008;**30**; 2734-2753.
- [3] Bosetti V, Carraro C, Duval R, Tavoni M. What should we expect from innovation? A model-based assessment of the environmental and mitigation cost implications of climate-related R&D. *Energy Economics*. 2011;**33**;1313-1320.
- [4] Nordhaus W. Modeling induced innovation in climate change policy. In: Grubler AN, *Technological Change and the Environment*. Washington D.C.: Resources for the Future Press. 2002
- [5] McJeon H, Clarke L, Kyle P, Wise M, Hackbarth A, Bryant B, et al. Technology interactions among low-carbon energy technologies: What can we learn from a large number of scenarios? *Energy Economics* 2011;**33**; 619-631.
- [6] Webster M, Paltsev S, Parsons J, Reilly J, Jacoby J, Uncertainty in Greenhouse Emissions and Costs of Atmospheric Stabilization. *MIT Joint Program Report 165*. Cambridge. 2008
- [7] Popp D. ENTICE: endogenous technological change in the DICE model of global warming. *Journal of Environmental Economics and Management* 2004;**48**;742-768.
- [8] McDonald A, Schramm L. Learning rates for energy technologies. *Energy Policy* 2001;**29**, 255-261
- [9] Rubin E, Taylor M, Yeh S, Hounshell D. Learning Curves for Environmental Technology and Their Importance for Climate Policy Analysis. *Energy* 2004;**29**;1551-1559.
- [10] Ringuest J, Graves S, Case R. Mean–Gini analysis in R&D portfolio selection. *European Journal of Operational Research* 2004;**154**;157-169.
- [11] Markowitz H. *Portfolio Selection: Efficient Diversification of Investments*. New York: Wiley; 1959
- [12] Samuelson P. Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics* 1969;**51**;239-247.
- [13] Chien CF. A portfolio evaluation framework for selecting R&D projects. *R&D Management* 2002;**32**;359– 368.
- [14] Avramov D, Zhou G. Bayesian Portfolio Analysis. *Annual Review of Financial Economics* 2010;**2**;25-47.
- [15] Zellner A, Chetty K. Prediction and decision problems in regression models from the bayesian point of view. *Journal of the American Statistical Association* 1965;**60**;608-616

- [16] Vilkkumaa E, Liesiö J, Salo A. Bayesian evaluation and selection strategies in portfolio decision analysis. Preprint submitted to the *European Journal of Operational Research*. September 12th, 2012
- [17] Chao HP, Peck S, Wan S. A strategic R&D planning model with an application to an energy research organization. *Applied Stochastic Models and Data Analysis* 1990;6:215-235.
- [18] Guo L. A Bayesian approach to energy R&D portfolio analysis. University of California Santa Barbara; 2011.
- [19] Flyvbjerg B, Skamris B, Buhl S. How common and how large are cost overruns in transport infrastructure projects? *Transport Reviews* 2003;23:71-88
- [20] IEA. *Cost and performance of carbon dioxide capture and storage from power generation*. International Energy Agency. Paris: IEA. 2011
- [21] Zero Emissions Platform. The Costs of CO2 Capture, Transport and Storage - Post-demonstration CCS in the EU. 2011.
- [22] Klif. *Klimakur 2020 (In Norwegian)*, Norwegian Climate and Pollution agency, Ch 13. The capture, transport and storage of CO2. 2011
- [23] Alstom. *Cost assessment of fossil power plants equipped with CCS under typical scenarios*. Paris: Alstom. 2011
- [24] IEA. CO2 capture and storage: A Key Abatement Option. International Energy Agency. Paris, France. 2008
- [25] Department of Energy (DOE). *FY 2013 Congressional Budget Request*. U.S. Department of Energy. 2012

Appendix A, Deriving analytical learning expression

Given two Gaussian functions g and f , where f is the pdf of cost and g is the pdf of the mean of f , we have:

$$g(\mu|m_t) = \frac{1}{s_t\sqrt{2\pi}} e^{-\frac{1}{2s_t^2}(\mu-m_t)^2}$$

$$f(c|\mu) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2\sigma^2}(c-\mu)^2}$$

We now seek to use a realized cost \hat{C} to update our parameters m and s for g . Bayes theorem gives us

$$g(\mu|c) = \frac{g(\mu)f(c|\mu)}{\int g(\mu)f(c|\mu)d\mu}$$

Using equation (1) and (2) and ignoring any constants yields

$$g(\mu)f(c|\mu) = e^{-\frac{1}{2}\left[\frac{(\mu-m_t)^2}{s_t^2} + \frac{(c-\mu)^2}{\sigma^2}\right]} = e^{-\frac{1}{2\sigma^2 s_t^2 / (\sigma^2 + s_t^2)}\left[\mu - \frac{\sigma^2 m_t + s_t^2 c}{\sigma^2 + s_t^2}\right]^2}$$

In other words we have

$$g(\mu|c) \propto N(m_{t+1}, s_{t+1}^2)$$

With

$$s_{t+1} = \sqrt{\frac{\sigma^2 s_t^2}{\sigma^2 + s_t^2}}$$

$$m_{t+1} = \frac{\sigma^2 m_t + s_t^2 \hat{C}}{\sigma^2 + s_t^2}$$

Starting with an initial guess of the parameters of g , and knowing σ , we now have an analytical expression for updating (m, s) for each cost \hat{C} that is observed.

Assuming that probability distributions for capture and storage costs are independent and identically distributed, the mean and standard deviation of the convoluted total cost function will be

$$s_T = \sqrt{s_{c,T}^2 + s_{s,T}^2}$$

$$m_{total} = m_c + m_s$$