



Should We Give Up After Solyndra? Optimal Technology R&D Portfolios under Uncertainty

Mort Webster
Karen Fisher-Vanden
David Popp
Nidhi Santen

CESIFO WORKING PAPER NO. 5448
CATEGORY 10: ENERGY AND CLIMATE ECONOMICS
JULY 2015

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

ISSN 2364-1428

Should We Give Up After Solyndra? Optimal Technology R&D Portfolios under Uncertainty

Abstract

Global climate change and other environmental challenges require the development of new energy technologies with lower emissions. In the near-term, R&D investments, either by government or the private sector, can bring down the costs of these lower emission technologies. However, the results of R&D are uncertain, and there are many potential technologies that may turn out to play an effective role in the future energy mix. In this paper, we address the problem of allocating R&D across technologies under uncertainty. Specifically, given two technologies, one with lower costs at present, but the other with greater uncertainty in the returns to R&D, how should one allocate the R&D budget? We develop a multi-stage stochastic dynamic programming version of an integrated assessment model of climate and economy that represents endogenous technological change through R&D decisions for two substitutable non-carbon backstop technologies. Using the model, we demonstrate that near-term R&D into the higher cost technology is justified, and that the amount of R&D into the high cost technology increases with both the variance in the uncertainty in returns to R&D and with the skewness of the uncertainty. We also present an illustrative case study of wind and solar photovoltaic technologies, and show that poor R&D results in early periods do not necessarily mean that investment should not continue.

JEL-Code: O380, Q420, Q480, Q550.

Keywords: innovation, climate change, solar energy, wind energy.

Mort Webster

*Energy and Mineral Engineering
Pennsylvania State University
USA - State College, PA 16802
mort@psu.edu*

Karen Fisher-Vanden

*Agricultural Economics, Sociology, and
Education / Pennsylvania State University
USA - University Park, PA 16802
kaf26@psu.edu*

David Popp

*Department of Public Administration and
International Affairs / Syracuse University
USA – Syracuse, NY 13244-1020
dcpopp@maxwell.syr.edu*

Nidhi Santen

*Synapse Energy Economics, Inc.
USA - Cambridge, MA 02139
nrsanten@synapse-energy.com*

1. Introduction

A response to global climate change will necessarily include reducing the carbon emissions from energy production in the future. There are many potential fuels and technologies that may contribute to future energy production with lower greenhouse gas emissions. For electric power, candidates include nuclear power, carbon capture and sequestration from coal or natural gas combustion, solar power (photovoltaic (PV) or concentrated solar thermal), wind power (on-shore or off-shore), and biomass-fueled combustion. Similarly, for transportation lower carbon technologies include efficient diesel engines, electric vehicles, hybrid electric vehicles, compressed natural gas (CNG) vehicles, and biofuels.

Potential low carbon technologies vary widely in their current relative costs, technological maturity, and commercial viability. One current policy challenge is how to efficiently allocate scarce resources for research and development (R&D) across different technologies. Public and private entities that fund R&D must choose how to allocate their efforts among these disparate technologies. For example, in the U.S., the Department of Energy (DOE) is a major source of public funding for energy R&D (DOE, 2015). Each year, DOE allocates funding for R&D across technology areas as part of their input into the budget and appropriations process. However, there is limited objective or quantitative guidance for agencies on how to choose among possible allocations (Pugh et al., 2011).

To focus the discussion, consider the simple case of two low-carbon energy technologies that are relatively easily substitutable. Although no technology is a perfect substitute for another, the ease of substitution in some circumstances may be significant. For example, solar PV and on-shore wind power constitute one such pair within the electric power system. Both are intermittent renewable sources of electricity that emit no carbon. Although their relative

potential in any specific location may vary depending on the solar and wind resource that exists, they play a similar role within the power system. Other possible examples include nuclear power and coal-fired electricity with carbon capture for electricity production (both are baseload technologies), or hybrid, battery electric, and CNG vehicles for transportation.

In each of these examples, one technology is currently significantly lower cost than the alternative(s). For example, the 2015 estimates of levelized cost from the U.S. Energy Information Administration is \$73.6 per MWh (2012 U.S. \$) for on-shore wind as compared to \$125.3 per MWh for solar PV (EIA, 2015). In this case, solar PV is still higher cost than wind despite decades of significant public and private R&D in many industrialized countries into solar PV technology, which has lowered its costs by several orders of magnitude (Barbose et al., 2014). A reasonable question, therefore, is whether continued R&D investments in solar PV are justified and at what level of funding, given that wind is lower cost and can fill a similar niche. This question, as applied to solar, has become particularly relevant since the high-profile bankruptcies of two solar technology companies in the U.S. that had received substantial public investment. Solyndra, a California-based company that manufactured solar panels, received \$539 million in guaranteed government loans in 2009, but by August 2011 had gone bankrupt (Weiner, 2012). Around the same time, Evergreen Solar, a Massachusetts-based solar panel manufacturer also declared bankruptcy (Church, 2011). In both cases, the reasons for the firms' failures were due to a number of factors, but certainly include competition from lower-cost manufacturers in China. Nevertheless, these failures have been cited in debates over public R&D, and some continue to argue that further investments in solar technology are not cost-effective (e.g., U.S. News & World Report, 2012), or at least that public investment in solar should be scaled back.

In this motivating example of wind vs solar PV, a critical factor is that uncertainty in the returns from R&D into solar is likely to be highly skewed, and more skewed than for wind, the more commercially established technology. For example, Baker et al. (2009a) performed an expert elicitation of the uncertainty in the results of solar R&D with several technology experts, and the resulting distributions are in fact highly skewed. In the past decade, numerous potentially revolutionary technologies for PV have been identified (e.g., Green, 2000) and much R&D invested by programs such as Sunshot within the U.S. Department of Energy (Sunshot, 2012), including advanced materials and manufacturing processes. In the past several years, prices have in fact declined sharply, but the cost reductions are primarily due to other factors such as economies of scale in production rather than new advanced materials (Barbose et al., 2014). Does this mean that these recent R&D investments were not wise, given the information available, or that continued investments are not warranted? Or do the results of solar R&D to date merely reflect a highly skewed distribution of uncertain returns, and is continued R&D still optimal?

In this paper, we investigate the factors that influence the economically efficient allocation of R&D investment across two substitutable energy technologies. In particular, we focus on the question: under what circumstances should we allocate some R&D to a higher-cost technology, when a lower cost substitute exists? Further, how should the allocation change if R&D into the higher-cost technology has uncertain and positively skewed returns? Here, we explicitly consider R&D as a stochastic process. Uncertainty is a central feature of R&D processes, since a priori one does not know what will be obtained from a given level of investment. Moreover, the relevant features of the uncertainty include more than simply mean and variance. Some technologies may have very low probability of a significant cost reduction

and very high probability of little or no cost reduction; this implies skewness in the distribution. When decision-makers must allocate R&D among technologies with qualitatively different risk profiles (e.g., more skewed but lower mean vs. less skewed and higher mean), how should this allocation be made? Using our stochastic framework, we demonstrate the influence of the shape of the probability distribution of uncertain returns to R&D on the optimal level of R&D investment into the higher-cost technology.

Numerous studies exist that have addressed the empirical, theoretical, and applied numerical modeling of innovation and the R&D process. Several excellent reviews exist (e.g., Pizer and Popp, 2008; Gillingham et al., 2008; Loschel, 2002, Popp, Newell, and Jaffe, 2010), so we only briefly review the most relevant studies here. Many theoretical and numerical models assume an innovation possibilities frontier (IPF) (Kennedy, 1964; Kamien and Schwartz, 1968; Acemoglu, 2002), which is a function of the stock of “knowledge capital” (Romer, 1990; Acemoglu, 1998). Examples of numerical models that include endogenous directed technological change include Popp (2004), Popp (2006), Buonanno et al. (2003), Sue Wing (2006), Goulder and Schneider (1999), and Goulder and Mathai (2000). The majority of these models are deterministic, and focus on specific questions such as the welfare gain from including endogenous technological change when there is a climate policy target (e.g., Popp, 2004; Buonanno et al., 2003; Goulder and Mathai, 2000), the relative impacts of various policy instruments on inducing innovation (e.g., Fischer et al., 2013; Fischer and Newell, 2008), the impact on optimal R&D of uncertainty in the extent of climate change (e.g., Baker et al., 2006), and the effect on R&D of spillovers across countries (e.g., Bosetti et al., 2009).

Some studies have directly addressed the uncertainty in the returns to R&D, but have tended to focus on R&D into a single backstop that substitutes for a conventional (i.e., fossil)

energy source. Theoretical models of optimal R&D under uncertainty include Hung and Quyen (1993), who showed that uncertainty can delay R&D and accelerate conventional resource depletion under some conditions, and Tsur and Zemel (2005), who explored a range of dynamic growth paths with and without R&D and showed that the growth path depends on the type of economy. Goeschl and Perino (2009) present a model in which the backstop technology has uncertain environmental characteristics in a multi-pollutant context, but deterministic cost impacts of R&D. Baker and Adu-Bonnah (2008) develop a two-stage stochastic model in which the impacts of R&D are uncertain with a probability distribution over three possible cost outcomes (target cost, breakthrough, or failure). They show that the effect of increasing the mean-preserving variance on R&D investment is ambiguous, and depends on the type of technology targeted. Bosetti and Tavoni (2009) present a stochastic two-stage version of the WITCH model (Bosetti et al 2006) in which the investment cost of the backstop is drawn from a symmetric distribution that depends on the cumulative knowledge stock. In contrast to Baker and Adu-Bonnah (2008), they demonstrate from both an analytical model and from the numerical model that uncertainty in the returns to R&D unambiguously increases the optimal level of R&D.

There are fewer studies that explicitly consider multiple technologies and uncertainty. One notable example is Gritsevski and Nakicenovic (2000), who considered more than 100 energy technologies within the MESSAGE model (Messner et al., 1996). However, their approach is focused on learning-by-doing rather than explicit R&D. Using a large number of scenarios of technology dynamics, they show that there are many distinct paths with similar energy costs. Pugh et al (2011) also used a scenario-based approach to explore the energy R&D portfolio question. Using results from a range of scenarios from the GCAM model (Brenkert et

al., 2003) to inform the benefits of different technologies and a probability distribution for the likelihood of success of achieving the maximum emissions reductions, they develop example 5-year R&D plans.

Most similarly to this study, Baker and Solak (2011) develop a two-stage (stage 1 chooses R&D and stage 2 chooses abatement) stochastic model with R&D into three alternative backstop technologies: solar, coal with carbon capture, and nuclear. The probability distributions for the uncertain returns to R&D are taken from expert elicitations of the uncertainty in cost reductions from R&D into each of these technologies (Baker et al., 2008, 2009a, 2009b). In general, they find that the R&D portfolio is diverse across technologies and is robust to the uncertainty in climate damages. However, they do not investigate the sensitivity to different shapes of the distributions of uncertain returns to R&D, nor do they consider more than two stages. Previous work (Webster et al., 2012) has shown that for some problems, multi-stage ($N > 2$) may lead to very different stage 1 optimal strategies compared with two-stage formulations.

Several questions and challenges remain unaddressed by the extant literature. First, given two technologies with uncertain returns to R&D, how should R&D be allocated? In particular, the relative impacts of mean, variance, and higher order moments in the distribution of uncertain returns on the R&D allocation decision have not been methodically explored. Second, in a multi-stage context with repeated opportunities for R&D and observed outcomes, what is the conditionally optimal R&D, given failures in previous stages to reduce cost? Finally, a challenge remains in calibrating models of technological change using the empirical literature in the treatment of the depreciation or decay rate of the knowledge stock. Representations of the Innovation Possibilities Frontier assume a depreciation rate, based on findings in the empirical

literature that knowledge decays over time in terms of its effectiveness at generating new knowledge. However, a deficiency in current formulations is that they therefore assume that the cost reductions achieved also decay, implying that costs increase over time without a minimum threshold of new knowledge every period.

We build on the previous studies of optimal R&D under uncertainty to meet a climate target by investigating the optimal ratio of R&D investment across two generic backstop technologies. We develop a multi-stage stochastic dynamic model, based on ENTICE-BR (Popp, 2006) with R&D-induced technological change by means of innovation possibilities frontiers and cumulative knowledge capital stocks. Using this model, we explore the factors that determine the relative magnitude of R&D investment that should be allocated to a higher cost backstop when a lower cost technology exists. In particular, we demonstrate the effect of varying the shape of the distribution of uncertain returns to R&D, and show that the optimal R&D depends not only on the relative variance but also on the skewness of the distributions. Our specific contributions are:

- A model of multi-stage R&D decisions across two low-carbon backstop technologies under uncertainty in returns to R&D;
- A novel formulation of technological change that distinguishes between knowledge stocks that generate new knowledge and those that generate perpetual cost reductions; and
- Demonstration of the relative effects of variance of the distribution of uncertain returns, skewness of the distribution, and initial relative costs of the backstops on the optimal R&D ratio.

The results of the analysis contribute to the policy discussion over R&D portfolios by demonstrating that for long-tailed (positively skewed) distributions, optimal R&D into the more

expensive backstop should be greater than for less skewed distributions, and that even below-average outcomes from first stage R&D into the high-cost technology does not mean that future R&D should be focused exclusively on the lower cost technology.

The rest of the paper proceeds as follows. In Section 2, we begin by presenting a simple analytical two-stage model of R&D allocation under uncertainty. This model frames the qualitative result that we later demonstrate with the numerical model. Section 3 presents the methodology and key assumptions for the stochastic dynamic version of ENTICE-BR with two substitutable backstops. We present the results from the numerical model in Section 4. Section 5 gives a concluding discussion.

2. R&D under Uncertainty and the effect of Skewness

To motivate and frame the numerical model and its results, we begin by illustrating the concepts with a simple stylized analytical model. Although this model omits many features of the problem to retain analytical tractability and transparency, we illustrate the key intuition. The numerical model presented in the subsequent sections then includes many of the complicated features omitted here, and produces qualitatively similar results.

Consider a two-stage model, $t = 1, 2$, and two possible backstop energy technologies, $i \in \{1, 2\}$. In the first stage ($t=1$), the decision maker must choose the R&D investment levels into each of the two technologies, R_1, R_2 . For simplicity, assume that the total budget is exogenously constrained:

$$R_1 + R_2 \leq B,$$

and that the two technologies are perfect substitutes.

We denote the cost of technology i at time t by C_i^t . Assume that Technology 1 is lower cost in the initial period than Technology 2:

$$C_1^1 < C_2^1$$

In this model, the cost in stage 2 can be reduced by R&D according to some increasing, concave function f :

$$C_i^2 = C_i^1 - f(R_i); \quad f' > 0, f'' < 0$$

The concavity assumption reflects the diminishing marginal returns from R&D into a given technology at a given time, which has been demonstrated empirically (e.g., Popp, 2002).

We model uncertainty in the returns to R&D into Technology 2 using a simple multiplicative shock to the cost reduction function. Specifically, we assume that with probability p , there is a multiplicative shock θ to $f(R_2)$, where $\theta > 1$. With probability $(1-p)$, the cost reduction in period 2 from R_2 is $f(R_2)$. This is a stylized approach, for simplicity and clarity, in which the probability p represents a realized cost reduction; in reality there is a sequence of conditions that must occur before the cost reduction occurs, including successful innovation, diffusion, etc., which we do not explicitly treat here. Furthermore, we assume that the magnitude of the exogenous cost shock $\theta > \theta^*$ is sufficiently large that it will make Technology 2 lower cost and therefore preferred to Technology 1 in the second period. For ease of explication, we neglect the case where $\theta < \theta^*$; this case has a trivial result since all R&D should go to Technology 1.

Because the motivation for this research is emissions reduction to address climate change, we assume that there is an emissions target in stage 2 that must be met. Because we have two non-carbon energy technologies that are perfect substitutes, the decision maker will adopt the technology that has the lowest costs in stage 2. In this case, we can frame the objective

function for the R&D decision as one of minimizing the expected cost of whichever technology is adopted. Specifically, we assume that with probability p , Technology 2 will be adopted with a cost of $C_2^1 - \theta f(R_2)$, and with probability $(1-p)$ Technology 1 is adopted with a cost of $C_1^1 - f(B - R_2)$. We therefore write the objective function as

$$\min_{R_1, R_2} p[C_2^1 - \theta f(R_2)] + (1 - p)[C_1^1 - f(B - R_2)]. \quad (1)$$

Taking the derivative and setting to zero, the first-order condition is:

$$-\theta p f'(R_2) + (1 - p) f'(B - R_2) = 0.$$

Rearranging this, we can express the optimality condition for R&D in Technology 2 as:

$$f'(R_2) = f'(B - R_2) \cdot \frac{1-p}{p\theta}. \quad (2)$$

The intuition for the expression in equation (2) is as follows. Relative to a solution where R&D is allocated such that the marginal productivities of R&D are equal (as would be implied by $f'(R_2) = f'(B - R_2)$), the effect of the second term on the RHS is that as the probability p of a better than expected shock increases, or as the magnitude of that shock θ increases, the optimal allocation will consist of R_2 at a lower marginal productivity than R_1 . From the concavity assumption, this in turn implies that R_2 will be increasing in both p and θ .

Put more simply, this result indicates that as the probability of a higher than average shock or the magnitude of a higher than average shock increases, increased R&D investment into the expensive technology will be optimal. As we will show below in the numerical results, this is particularly true for positively-skewed distributions. Because a probability distribution must integrate to unity, a necessary condition for increasing the magnitude of an extreme upper quantile while holding the expected value constant is that there must be correspondingly more probability mass below the mean of the distribution. In such a distribution, proportionally more random draws will yield below average returns. In order to achieve the same expected cost

reduction as a distribution with the same mean but less skewed, more random draws are required. One can think of “more draws” as roughly implying greater R&D investment. In this stylized example, greater skewness therefore implies that more R&D investment is optimal.

The model presented above is necessarily simplified and omits many critical features. Such features include the presence of cumulative, depreciating knowledge capital stock, a non-linear formulation for the innovation possibilities frontier, a dependence of the benefits of R&D on the exact amount of future adoption of each technology, less than perfect substitutability, and multiple R&D decision stages. Because simple analytical closed-form solutions are not possible for a model with these features, we turn for the remainder of the paper to a numerical model.

3. Numerical Model Methodology

For a numerical example, we build on the ENTICE (Popp, 2004) and ENTICE-BR models (Popp, 2006a, 2006b), both of which are extensions of the DICE model (Nordhaus and Boyer, 2000). ENTICE and ENTICE-BR are integrated assessment models of climate change that represent the global economy and environment and include endogenous technological change. ENTICE included R&D into energy efficiency among the decision variables, and ENTICE-BR added R&D into a non-carbon emitting energy backstop technology.

We modify the ENTICE-BR model to address the question of what is the optimal R&D allocation across two substitutable energy backstops under uncertainty. The original ENTICE-BR model includes R&D investment to augment two knowledge capital stocks according to an innovation possibilities frontier, one which increases the energy efficiency of the economy and one which lowers the cost of a non-carbon energy backstop. The original model is deterministic,

and the parameters in the model were calibrated to be consistent with historically observed patterns of R&D, energy patents, and energy consumption (Popp, 2001; Popp, 2002).

Here, we introduce three major modifications to the original ENTICE-BR model: 1) we replace the single non-carbon energy backstop with two backstops that can substitute for each other; 2) we modify the representation of the knowledge stock to distinguish the dual roles for knowledge capital of generating new knowledge and reducing technology cost; and 3) the model is reframed as a stochastic dynamic programming problem, and is solved using approximate dynamic programming. Finally, for the case study of wind vs. solar in the following section, we calibrate parameters in this model to be consistent with results from Popp et al. (2013), which used forward patent citation counts by energy technology type to explore the distribution of uncertain returns to R&D and the decay rate of knowledge capital by technology. Each of the above modifications is described in detail below.

3.1 Two Backstop Technologies

The DICE/ENTICE family of models all have the objective of maximizing the present value of discounted utility, subject to economic and climate system constraints. Here we discuss only the relevant constraints that are modified, and provide the full model description in the Appendix. In the original ENTICE-BR, output Q_t is produced by physical capital stock K_t , labor L_t , and effective energy units E_t . Labor grows according to an exogenous population growth trend, and physical capital stock is the result of investment plus the previous period's stock that has not depreciated. Overall technological progress is modeled with total factor productivity A_t . Effective energy units are produced from three possible energy inputs: fossil fuels F_t , a carbon-free backstop B_t , and knowledge pertaining to energy efficiency $H_{E,t}$. The cost of fossil fuels and

backstop energy, respectively are $p_{F,t}$ and $p_{B,t}$, and are subtracted from total output. Thus the production function for the economy is:

$$Q_t = A_t K_t^\gamma L_t^{1-\gamma-\beta} E_t^\beta - p_{F,t} F_t - p_{B,t} B_t. \quad (1)$$

Effective energy units are produced from a nested constant elasticity of substitution (CES) representation that aggregate fossil fuels F_t , backstop energy B_t , and knowledge stock for energy efficiency $H_{E,t}$:

$$E_t = \left[\alpha_H H_{E,t}^{\rho_H} + \left(\left(\frac{F_t}{\alpha_\phi \Phi_t} \right)^{\rho_B} + B_t^{\rho_B} \right)^{\rho_H/\rho_B} \right]^{1/\rho_H}. \quad (2)$$

The elasticity of substitution between fossil fuels and backstop energy is $1/(1 - \rho_B)$, and the elasticity of substitution between the fossil/backstop bundle and energy efficiency is $1/(1 - \rho_H)$.

We extend this model to include two carbon-free energy backstops, $B_{1,t}$ and $B_{2,t}$. Thus equation (2) is replaced with:

$$E_t = \left[\alpha_H H_{E,t}^{\rho_H} + \left(\left(\frac{F_t}{\alpha_\phi \Phi_t} \right)^{\rho_B} + (B_{1,t}^{\rho_\kappa} + B_{2,t}^{\rho_\kappa})^{\rho_B/\rho_\kappa} \right)^{\rho_H/\rho_B} \right]^{1/\rho_H}. \quad (3)$$

The elasticity of substitution between the two backstops is $1/(1 - \rho_\kappa)$. Similarly, we denote the respective costs of the backstops with $p_{B1,t}$ and $p_{B2,t}$ and subtract these from total output.

Equation (1) is therefore replaced by:

$$Q_t = A_t K_t^\gamma L_t^{1-\gamma-\beta} E_t^\beta - p_{F,t} F_t - p_{B1,t} B_{1,t} - p_{B2,t} B_{2,t}. \quad (4)$$

3.2 Knowledge Stock Representation of Technological Change

In ENTICE-BR, the cost of the backstop can be reduced by investing in a knowledge stock $H_{B,t}$. The knowledge stock increases with the accumulation of research and development $R_{B,t}$, as:

$$H_{B,t} = h(R_{B,t}) + (1 - \delta_H)H_{B,t-1} \quad (5)$$

The parameter δ_H represents knowledge decay. The function $h(R_{B,t})$ is the innovation possibility frontier, which models the process of new knowledge creation from R&D and the existing knowledge stock:

$$h(R_{B,t}) = \alpha R_{B,t}^{b_B} H_{B,t}^{\phi_B}. \quad (6)$$

Both b_B and ϕ_B are assumed to be between 0 and 1, and $\phi_B < 1$ represents diminishing returns to research across time (Popp, 2004). Finally, the relationship between accumulating knowledge stock and the cost of the energy backstop is represented by:

$$p_{B,t} = \frac{p_{B,0}}{H_{B,t}^\eta}. \quad (7)$$

The parameter η converts knowledge stock level to the reduction from the backstop cost in the initial period, and is calibrated so that $1 - 2^{-\eta}$ is the cost reduction from a doubling of the knowledge stock. Consistent with other treatments of endogenous technological change, we assume here that an increase in the knowledge stock necessarily leads to a reduction in the cost of the technology. In reality, new knowledge can sometimes increase cost, or even increase uncertainty, although such cases may be fairly nuanced, with accompanying increases in quality. For purposes of clarity, we do not include the possibility of increased cost from new knowledge here, but leave that to future work.

One of the challenges in calibrating representations of technological change such as described above using empirical data is that although there is evidence that knowledge decays

over time, using non-zero depreciation rates has the effect that a non-zero level of R&D is required to simply maintain the backstop cost at a constant level over time. Without new R&D investments in a model like equations (5-7), the cost of the backstop will rise over time as the knowledge stock decays. This behavior is not intuitive. Despite the evidence that knowledge's effectiveness at creating new knowledge decays, there is not comparable evidence that the reduction in the cost of a technology from the stock of knowledge decays. Once a cost reducing change to a technology has been innovated, diffused, and adopted, it remains part of the state of the art of that technology.

Here we introduce a modification to the prevailing representation of knowledge stock to address the challenge just described. Instead of a single knowledge stock for each technology, we separately track two distinct but related knowledge stocks for each technology. One knowledge stock, $H_{B,t}^K$, is combined with new R&D investment $H_{B,t}^K$ to create new knowledge (as in equation (6)). This first type of stock depreciates over time. The second stock, $H_{B,t}^C$, is used only for the reduction of technology costs (as in equation (7)). This latter stock does not depreciate over time, so that if R&D investments are set to zero after a given time, the technology costs will remain constant. This modeling innovation allows us to calibrate the model to fit historical observations on both wind and solar technologies for the case study presented in section 4.4. The dynamics and constraints for both stocks are defined for each of the two backstop technologies.

The model used here therefore includes the following constraints, which replace equations (5-7). The knowledge stock dynamics are:

$$H_{B,t}^K = h(R_{B,t}) + (1 - \delta_B)H_{B,t-1}^K \quad B \in \{B_1, B_2\}, \quad (8a)$$

$$H_{B,t}^C = h(R_{B,t}) + H_{B,t-1}^C. \quad (8b)$$

The innovation possibilities frontiers use only the knowledge creation stock:

$$h(R_{B,t}) = aR_{B,t}^{b_B}(H_{B,t}^K)^{\phi_B} \quad B \in \{B_1, B_2\}. \quad (9)$$

The cost of each backstop is a function of the cost-reduction knowledge stock:

$$p_{B,t} = \frac{P_{B,0}}{(H_{B,t}^C)^\eta} \quad B \in \{B_1, B_2\}. \quad (10)$$

3.3 Stochastic Dynamic Formulation

As described above, the objective function for the deterministic version of the model is to maximize the present value of discounted utility, subject to the economic and climate system constraints, or mathematically:

$$\max_{X(t)} \sum_{t=0}^T U(c(t), L(t))(1 + \rho(t))^{-1}. \quad (11)$$

$X(t)$ represents the decision vector over time, which includes investment in physical capital stock as well as R&D and the energy produced from fossil and the backstop sources. In general, we assume that utility is a function of per capita consumption $c(t)$ and population $L(t)$, and that $\rho(t)$ is the social rate of time preference.

To model decision under uncertainty, we reformulate the model as a stochastic dynamic program. We use a formulation consistent with that in Webster et al (2012), in which the DICE model was framed as a seven-stage stochastic dynamic program. Here, we also assume seven¹ decision stages for the reference version of the model $t = \{1, 2, \dots, 7\}$. The decision or action space of the model consists of the R&D investment into each of the two backstop energy technologies and the quantity of energy produced from each backstop in the current stage, $X =$

¹As shown in Webster et al. (2012) and also in the Appendix to this paper, the impact on the first stage decision of the number of decision stages is roughly the same for any number of stages T greater than five. The choice of seven stages is for computational tractability. Given sufficient computing power, the results for $T = 35$ (decadal decisions) would be qualitatively the same.

$[R_1, R_2, B_1, B_2]$. Uncertainty is represented in the returns to R&D as modeled by the innovation possibilities frontier (eq. 9). We introduce a multiplicative shock θ to the function $h(R_{B,t})$, and assume that θ is drawn i.i.d. from a distribution with an expected value of 1, such that the mean of many draws is equal to the original deterministic model. We consider many different distributions in the analysis, and describe them in the next section.

The solution of the finite horizon stochastic dynamic program described can be obtained by solving the Bellman equations (Bellman, 2003) for all states and all stages:

$$V_t = \max_{X_t} [U_t + E\{V_{t+1}(X_t, \theta_t)\}]. \quad (12)$$

The value function V_t is a mapping from state to optimal value, and the conventional solution method for a finite-horizon problem is backward induction. However, backward induction requires that the state space is defined such that the process is Markov. For the ENTICE-BR model including the modifications described above, this would require a twelve-dimensional state space (see Appendix). For a seven-stage model and any reasonable discretization resolution in each dimension of state space, the computational cost of exact DP would be prohibitive.

As a result, we solve the model using Approximate Dynamic Programming (ADP) (Bertsekas and Tsitsiklis 1996; Powell 2007). ADP methods find an approximate solution to the dynamic programming problem using adaptive random sampling to estimate an approximation of the value function, using basis functions defined on a subset of the dimensions of the state space. The specific algorithm is similar to that applied in Webster et al. (2012), but with a few key differences, described here; the full description may be found in the Appendix. In this study, we approximate the value function as a linear function of the prices of the two backstops:

$$\hat{V}_t = \beta_{0,t} + \beta_{1,t}p_{B1,t} + \beta_{2,t}p_{B2,t}$$

The vector β is solved for using least squares on an initial set of sample paths, and then is iteratively updated using the Bellman Error method (Bertsekas and Tsitsiklis 1996) until convergence.

4. Numerical Case Study: R&D Allocation Across Two Backstops

In this section, we explore the factors that determine optimal allocation of R&D across two substitutable energy backstops in the ENTICE-BR model. We begin with the dynamics when the returns to R&D are deterministic. We then show how the results change when the returns to R&D are stochastic, using a reference distribution for the shock to R&D returns. We then focus on a case where the cost in the initial period of Backstop 2 is 50% greater than that of Backstop 1, and show the sensitivity of optimal R&D for Backstop 2 to variance and skewness of the distribution. Finally, we use parameters and distributions calibrated to Popp et al (2013) as described above to simulate a case study of Wind vs Solar R&D.

4.1 Deterministic Results

In the absence of uncertainty, the major determinant of the optimal R&D shares is the ratio of the backstop prices. The benefit of R&D in each period is the reduction in expected future costs. Therefore, backstop energy use in future periods largely determines the R&D into that technology in each period. To illustrate this effect, we present the optimal backstop use and R&D for both technologies for three different initial prices of Backstop 2, \$600, \$1200, and \$1800 per ton of carbon equivalent (CTE) (Fig. 1). Based on Popp (2006a), conventional fossil energy is \$276 CTE, so the assumed initial backstop prices in the results below are slightly more than double, four times, and six times the cost of fossil energy. In all cases, the initial price of

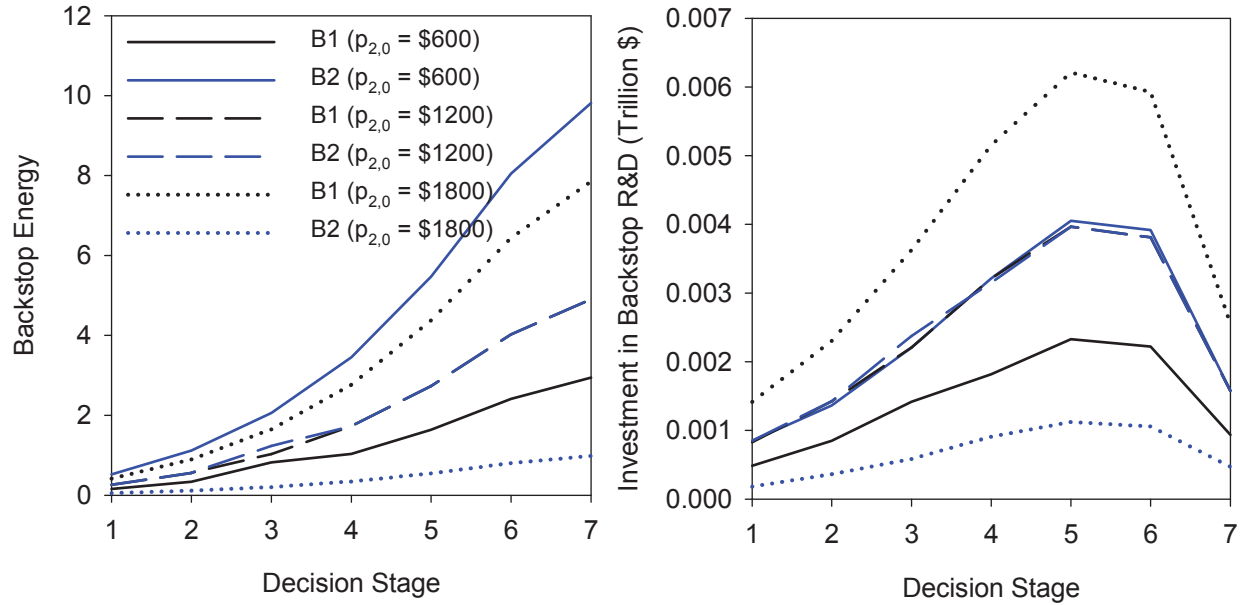


Figure 1: Energy produced from each backstop (left) and R&D investment into each backstop under initial prices for Backstop 2 of \$600, \$1200, and \$1800 (Backstop 1 is \$1200 in the initial period for all cases).

Backstop 1 is \$1200. The elasticity of substitution between the two backstops is assumed to be 5.0, which is highly elastic and makes the two backstops behave as nearly perfect substitutes.

As expected for substitutable technologies, when the initial prices are equal (both at \$1200), the backstop use and the R&D are equally divided between the two technologies. When the initial price of Backstop 2 is \$600, substantially more of this technology is used to provide non-carbon energy, and less of Backstop 1 is used. The R&D into Backstop 1 is also lower in this case. When the initial price of Backstop 2 is \$1800, more energy is produced from Backstop 1 and less from Backstop 2, and a similar pattern in R&D investment allocation is observed. Note that the low and high initial price cases are not symmetric, because the average cost over both backstops are not equivalent across cases. The total energy produced from both backstops is greatest for initial prices of \$1200 (B1) and \$600 (B2), and the total energy is the lowest for

Table 1: Total R&D and Backstop Energy under Deterministic Returns to R&D

Initial B2 Price		Total R&D Investment (Billions of 1995 US \$)			Total Energy from Backstops (tons of carbon equivalent)		
		600	1200	1800	600	1200	1800
Decision Stage	1	0.0013	0.0017	0.0016	0.7	0.5	0.5
	2	0.0022	0.0028	0.0027	1.5	1.1	1.0
	3	0.0036	0.0046	0.0042	2.9	2.3	1.9
	4	0.0050	0.0064	0.0061	4.5	3.5	3.1
	5	0.0064	0.0079	0.0073	7.1	5.5	4.9
	6	0.0061	0.0076	0.0070	10.5	8.0	7.2
	7	0.0025	0.0032	0.0030	12.8	9.8	8.8

initial prices of \$1200 (B1) and \$1800 (B2) (Table 1). Total R&D, however, follows a different pattern, with total R&D across both technologies being greatest for the case with initial prices of \$1200 (B1) and \$1200 (B2), and least total R&D is for the case where initial prices are \$1200 (B1) and \$600 (B2). When prices are lower for both backstops from the beginning, the marginal benefit of R&D (i.e., the avoided future cost of backstop energy) is lower, which leads to less R&D investment. Note that in this model we only consider cost reductions from R&D and omit any learning-by-doing effects that would yield additional cost reductions in the technologies as a function of cumulative production.

4.2 Stochastic Results for a Reference Distribution

To further investigate the dynamics of R&D allocation across the two backstops, we next introduce uncertainty in the returns to R&D for Backstop 2 (Backstop 1 R&D remains deterministic in this subsection). As described in Section 3, the quantity of new knowledge created from R&D into Backstop 2 (eq. 9) is multiplied by a random shock θ that is drawn from a Normal distribution with a mean of 1.0 and a standard deviation of 0.3. In the next subsection,

we will perform sensitivity analysis on the distribution of the shock. As above, we assume here that the initial price for Backstop 1 is \$1200, and that the elasticity of substitution between the backstops is 5.0 (highly elastic).

For all stochastic results shown here and in the rest of the paper, the solution method is as follows. First, the ADP algorithm is used to solve for the approximate value functions. After the value functions have converged, an additional set of random sample paths are simulated, with optimal decisions chosen on the basis of the value functions. For this set of sample paths, there is an optimal first stage decision and a probability distribution for latter stage decisions and other endogenous results. We will present most results in terms of the median and 90% probability range for all backstop energy and R&D investment decisions in Stages 2 through 7.

Consider first an initial cost of Backstop 2 of \$600, much lower cost than Backstop 1 with a cost of \$1200. For all sampled paths of new knowledge creation for Backstop 2, the energy used from the backstops is the same, consisting of slightly over 3 times as much from Backstop 2 as from Backstop 1 (Fig. 2a,b). In contrast, the paths of optimal R&D into backstop 2 do vary with each sample path, with a 90% probability range across the samples of +/- 25%. When better than average shocks to knowledge creation are observed, more R&D into Backstop 2 is optimal. The marginal benefit of additional R&D into Backstop 2 is a function of the expected avoided costs of using that backstop in all future periods. Because the backstop energy decisions are not changing for these parameter values, the observed variation in R&D is entirely due to changes in the productivity of R&D that results in larger or smaller cost reductions for future energy use; there is no substitution induced between the two backstops.




Figure 2: Optimal backstop energy and R&D investment for three initial prices.

Similarly, if the initial cost of Backstop 2 is \$1800, much higher than the cost of Backstop 1, the backstop energy decisions do not change across the sampled shocks to knowledge creation, but the optimal R&D into Backstop 2 does vary (Fig. 2e,f). In this case, roughly 8 times as much energy is produced from Backstop 1 as from Backstop 2. The relative variability in optimal R&D is similar to the \$600 case, growing to +/- 20%. However, the magnitude is much lower because less energy is produced from Backstop 2, and therefore the benefits of lowering its cost are reduced.

A different dynamic is observed for the case of equal initial costs of backstops (Fig. 2c,d). In this case, the energy produced from both backstops does vary with the particular sample path of shocks to knowledge creation. The median path of energy decisions is the same as the deterministic optimal path from the previous section, consisting of equal shares of energy produced from the backstops. However, better than average returns from R&D into Backstop 2 in early periods leads to greater energy production and greater R&D being optimal for Backstop 2, and correspondingly less from Backstop 1. The converse is also true for worse than average returns to R&D into backstop 2 in early periods. When technology costs are similar, the variation in optimal R&D reflects the aggregation of two effects: the benefit of lowering the cost of energy that would be used anyway (as in the \$600 and \$1800 cases), and additionally the benefit of lowering the cost of the technology sufficiently that it substitutes for the other higher-cost technology, which further reduces future total energy costs.

4.3 Sensitivity to Distribution Shape

We next vary the probability distribution from which the shock to knowledge creation is drawn. The purpose of this sensitivity analysis is to investigate the influence of characteristics of

the probability distribution on the optimal R&D portfolio across the two backstops. We consider a range of parameters for three different distribution families, the Normal or Gaussian distribution, the Lognormal distribution, and the Generalized Pareto distribution. The latter two are positively-skewed distributions, and the magnitude of the skewness varies with the parameters²; see Fig. 3 for an example probability density function from each of these families for one choice of distribution parameters. In all cases, the distribution for the multiplicative shock is normalized to ensure that the expected value is equal to 1.0, so that the mean of a large number of samples of new knowledge created by R&D will approximately equal the deterministic value. For the Lognormal distribution, we use parameters $\mu=3.0$ and vary σ over the range 0.5 to 1.5. For the Generalized Pareto distribution, we set the location parameter μ and the scale parameter σ both to 0.5, and vary the shape parameter ξ from 0.4 to 1.0.

We compare the optimal ratio of R&D across the two backstops for a range of initial costs of Backstop 2 (Backstop 1 remains \$1200 in the first period in all cases), and for several alternative assumptions about the uncertainty in returns to R&D (Fig. 4). The overall trend is dominated by the relative costs of the two technologies, independent of the representation of uncertainty. For the deterministic case (solid line in Fig. 4), the R&D into backstop 2 exceeds the R&D in Backstop 1 when the initial cost of Backstop 2 is below \$1200. Note that the impact of initial cost on R&D allocation is non-monotonic below \$1200. This behavior is a consequence of the interaction of the two effects demonstrated above. If the cost of Backstop 2 is moderately lower than the cost of Backstop 1 (e.g., \$800), then the benefits of R&D into Backstop 2 are the sum of 1) the benefits of reducing the cost of Backstop 2 energy that would

² Empirical research on the returns to R&D consistently finds evidence of highly skewed results -- some research projects result in grand successes worth billions of dollars, while others end in failure, with millions spent without any positive results. See, for example, Pakes (1986), Shankerman and Pakes (1986), Schanerkam (1998), and Jaffe and Trajtenberg (2002).

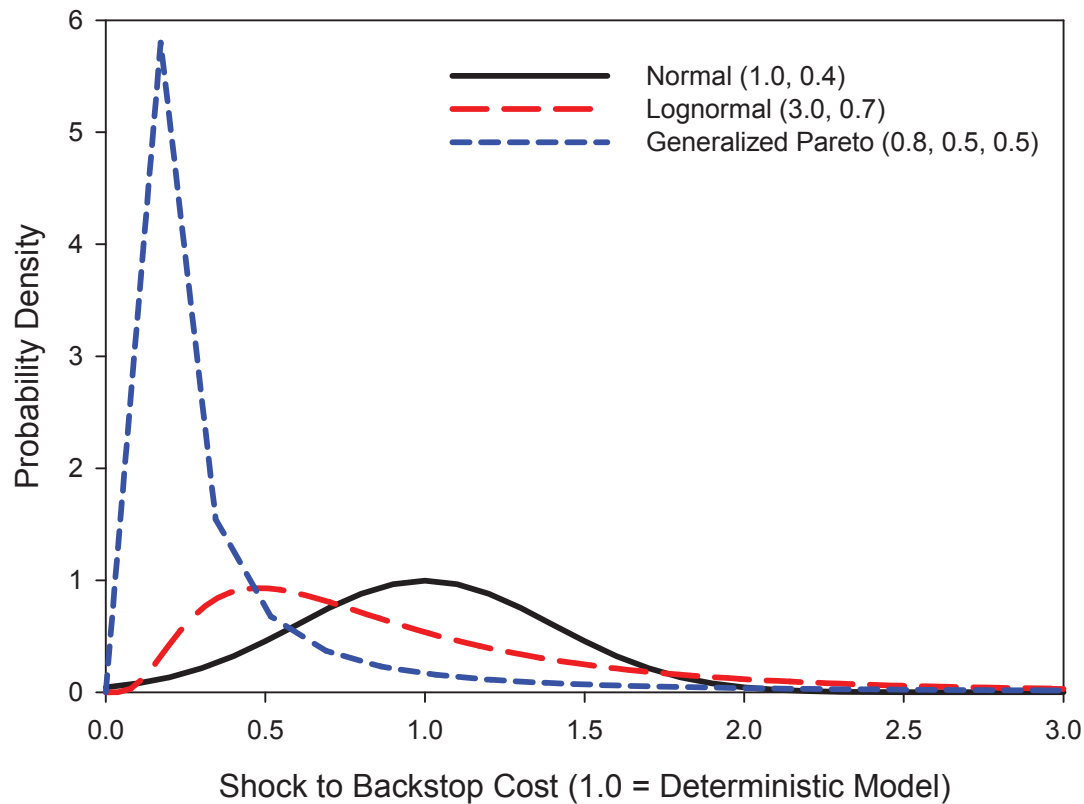


Figure 3: Example probability distributions. Shows the probability distributions of random shocks to backstop cost from which the model draws, with one example of Normal (solid black line), Lognormal (dashed red line), and Generalized Pareto (dashed blue line). Note that all distributions have mean value of 1.0.

have been used in the absence of further R&D, and 2) the benefit of reducing the cost of Backstop 2 sufficiently to substitute away from the higher-cost Backstop 1 and further reduce total energy costs. For very low initial costs (e.g., \$600), no substitution will occur and the cost before R&D is already low, so the marginal benefits of further cost reductions are smaller. When the initial costs of Backstop 2 are greater than \$1200, the optimal R&D for Backstop 2 is less than the R&D for Backstop 1, and is decreasing in the initial backstop cost.




Figure 4: Optimal R&D Ratio (Backstop 2 / Backstop 1) in Stage 1, as a function of initial cost of Backstop 2. The solid line indicates the optimal ratio with no uncertainty, the long dashed line indicates the ratio when the uncertainty in returns to R&D follows a Normal distribution ($\sigma=0.3$), the short dashed line is for the Lognormal distribution ($\sigma=0.7$), and the dotted lines is for a Generalized Pareto distribution with shape parameter $\xi = 0.4$.

Relative to the optimal R&D ratio in the deterministic case, the effect of uncertainty in the returns to R&D on Stage 1 is to increase the relative R&D investment in Backstop 2, relative to Backstop 1, for any level of initial backstop cost. We illustrate the effect of uncertainty on the optimal Stage 1 R&D ratio with three example distributions: Normal ($\sigma=0.3$), Lognormal ($\sigma=0.7$), and Generalized Pareto ($\xi=0.4$). Note that the relative order of optimal R&D ratio by distribution is not consistent over all initial costs of Backstop 2.

To further investigate the impact of the distribution shape on the optimal R&D portfolio, we focus on a relevant, if stylized, case of initial Backstop 2 cost at \$1800, 50% greater than the initial cost of Backstop 1. In a dynamic programming setting, the optimal decision is the one that maximizes the sum of the expected utility of the current stage and the expected value of being in a particular state at the next stage. For the R&D decision, the opportunity cost of the R&D causes the utility of the current stage to be decreasing in R&D. The expected value of the next stage depends on the expected costs of the backstops, and therefore is generally increasing in R&D investment. The sum of the two components is a convex function, as shown in Figure 5. Here, we plot the expected value of utility observed in stage 1 for different levels of R&D investment in backstop #2. The optimal R&D for any distribution occurs where this function is maximized, as shown by the vertical lines in Figure 5. The effect of an increase in uncertainty in returns to R&D is to increase the expected marginal utility of a unit of R&D in the current stage. The reason for this effect is the asymmetry in the impact of cost reductions. A smaller than average return to R&D produces little cost reduction, but the cost is not more than the current cost. A larger than average return to R&D produces a very large cost reduction and therefore significantly increases utility. The average over many smaller and larger than expected returns to R&D is a higher expected benefit of R&D, while the opportunity cost of R&D is unchanged. The net effect of increasing uncertainty in returns to R&D is to increase the optimal R&D investment into Backstop 2. A positively-skewed distribution (Lognormal or Pareto) results in an even higher optimal level of R&D (note that in Figure 5, the maximum of the curves are increasing as a more skewed distribution is assumed), suggesting that skewness may be a relevant factor in addition to variance. Another way to view the same result is to examine the values of the coefficients for the approximate value function for Stage 2 after convergence. The

approximate value function is a linear function of the costs of the two backstops. Note that the primary effect of varying the distribution of uncertain returns to R&D is to increase the magnitude of the coefficient for the cost of Backstop 2 (Table II). In other words, for higher variance and higher skewness, there is greater expected utility for a one dollar reduction in the cost of Backstop 2.

To further illuminate the influence of alternative distributions, we also show the median and 90% ranges for the optimal energy and R&D decisions over all Stages when the initial cost of Backstop 2 is \$1800 and the distribution of uncertain returns is either Lognormal or Pareto (Fig. 6). The relevant comparisons for these figures are with Figs. 2e and 2f, which show the equivalent results for a Normal distribution. Recall that with a Normally distributed shock and

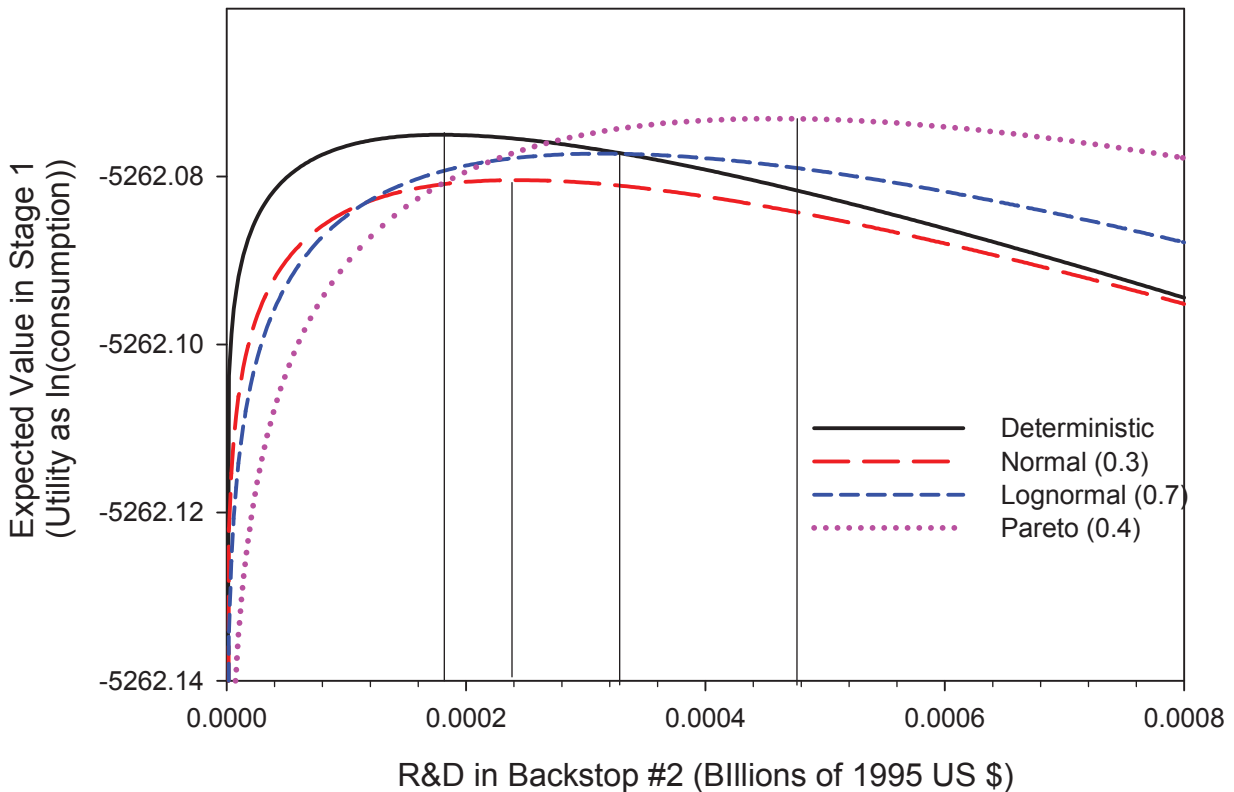


Figure 5: Expected value (utility) of Stage 1 R&D into Backstop 2 for alternative assumptions about uncertainty in returns to R&D. Initial cost of Backstop 2 is \$1800.

Table II: Final Coefficients for Approximate Value Function for Stage 2

Case	Constant	Slope for Price of B1	Slope for Price of B2
Deterministic	-1778.1	-0.0056	-0.00049
Normal (0.3)	-1777.7	-0.0055	-0.00078
Lognormal (0.7)	-1777.0	-0.0058	-0.00105
Pareto (0.4)	-1776.1	-0.0059	-0.00154




Figure 6: Optimal Backstop Energy Production and R&D for the Lognormal (0.7) and Pareto (0.4) distributions.

an initial cost of \$1800, the energy produced from Backstop 2 did not vary across the sample paths. In contrast, energy from Backstop 2 may increase under the Lognormal distribution, and the potential (upper 90%) energy production from Backstop 2 under the Pareto increases even further. Substitution of Backstop 2 for Backstop 1 may occur under these distributions, because of the long upper tail and the low probability of very large shocks to knowledge creation. As a consequence of this potential substitution, as well as the potential cost reduction, the corresponding variation in optimal R&D is also much larger than for the Normal distribution.

To characterize more generally the effect of the distribution shape on optimal R&D, we consider a larger set of alternative distributions for the returns to R&D on Backstop 2. Specifically, we solve for 17 distinct distributions, varying the parameters for Normal, Lognormal, and Generalized Pareto (indicated as black dots in Fig. 7). As above, we only consider an initial cost of \$1800 for Backstop 2. In general, two effects are apparent across these solutions. First, as the variance of the distribution increases, the increase in optimal R&D into Backstop 2 relative to the deterministic solution increases. In addition, an increase in the skewness of the distribution also increases the optimal R&D into Backstop 2. Plotting these effects against traditional moments of the distributions, variance and skewness (not shown), exhibits highly non-linear behavior. However, a transformation of variables, using the natural log of variance and the difference between 1.0 and the median of the distribution as a measure of skewness, display a fairly linear relationship between these variables and the relative increase in optimal stage 1 R&D (surface in Fig. 7). Note that by design, we have restricted all admissible distributions for this problem to ones with an expected value of 1.0. By definition, therefore, any positively-skewed distribution will necessarily have a median that is less than 1.0.




Figure 7: Optimal Stage 1 R&D into Backstop 2 relative to Deterministic Optimal R&D as a function of the log of the variance and the distance between 1.0 and the median of the distribution (as a measure of skewness). Dots indicate results for a specific distribution (i.e., given family and given parameter values), and the surface plot is the linear regression over all distributions simulated.

Recall that in Section 2, we derived an expression for the optimality conditions for the simple stylized model which showed that optimal R&D into the expensive backstop with uncertain returns was increasing in p , the probability of a higher than average shock, and increasing in θ , the magnitude of the higher than average return to R&D (eq. 2). An increase in p or in θ in this simple model is equivalent to increasing both the variance and the skewness of

the distribution. Thus, the general behavior of the numerical model is consistent with that of the simple analytical model.

4.4 Case Study: Wind vs. Solar

Here, we use the numerical model to examine R&D for a pair of technologies that are part of the public debate over low-carbon energy, wind and solar. These technologies are often compared because they are both non-carbon electricity sources that are intermittent. Typical estimates are that solar photovoltaic has 60% greater cost than on-shore wind. As discussed in the Introduction, some have argued based on recent “failures” in solar that further investment is not justified. We calibrate relevant parameters of our two-backstop model to approximate critical characteristics of wind and solar technologies. We then use the numerical model to explore optimal investment strategies in latter stages after observing below average R&D outcomes in previous stages.

For parameter calibration, we use results from Popp et al (2013), which characterized differences in patent forward citations by energy technology. By showing the links between past and current innovation, patent citations are an indicator of the social value of innovation as a building block for future research, and thus analogous to the variable $H_{B,t}^K$ in our model (Lanjouw & Schankerman 2004, Popp, 2002). One of the features shown to vary widely by technology was the effect of citation lag on the probability of citation. One can consider this as proxy for the decay rate of the knowledge stock used to create new knowledge for a given technology. The faster the probability of citation decreases with time, the faster the depreciation rate for the knowledge stock. We therefore calibrate our depreciation parameters for wind and solar to these results (see Fig. 5 in Popp et al., 2013). This procedure yields a depreciation rate

of 0.01 for wind and of 0.05 for solar, leading to faster knowledge decay in the latter. These estimates reflect the fact that wind and solar technologies are in very different stages of the innovation process, and wind is much less sensitive to knowledge stock decay as explained in more detail in Popp et al. (2013).

In terms of initial prices of the backstops, we keep wind (Backstop 1) at \$1200, and solar PV (Backstop 2) is assumed to be \$2000. This is based on the ratio of the current levelized costs of electricity estimated by the Energy Information Administration (EIA, 2015), which estimates that the levelized cost of solar PV is 67% greater than the cost of (on-shore) wind.

Modifying the depreciation rates alone leads to results that are inconsistent with the calibration of the rest of the model. We therefore repeat the calibration exercises of Popp (2006a), and adjust the scale parameter a in the IPF (see eq. 6) to ensure that the backstop cost trend over time is consistent with the costs produced by the original reference parameter values. We use the original single-backstop version of ENTICE-BR for the calibration exercise, and adjust parameter values for only wind and for only solar, respectively, applying each of the above depreciation rates. This procedure leads to values for a of 0.017 for wind and 0.0314 for solar. In other words, to preserve the price trend observed over time with faster depreciation requires correspondingly greater average effectiveness of R&D at creating new knowledge.

Finally, we calibrate the shape of the distribution of uncertain returns to R&D to match the shape of the distribution of patent forward citations for wind and for solar (see Fig. 2 in Popp et al., 2013). Using 1995 as a representative year, and normalizing so that the expected value of both distributions are 1, we find that for wind a Normal distribution with variance 0.5 and for solar an Exponential distribution with parameter $\lambda = 1$ best fit the shape of forward citations for the respective technologies.

To illustrate the results, we focus here on the ratio of energy production and the ratio of R&D investment between the two backstops (Solar / Wind), rather than the absolute magnitudes of R&D and energy decisions. Presenting the outcome in this form relates more directly to the objective of informing relative R&D allocation. We therefore explicitly report the ratio of backstop energy and R&D for each sample path after convergence, and summarize with the median and 90% range of that ratio over time.

The results are shown in Figure 8. If there were no uncertainty in returns to R&D for either technology, the energy production from Solar would be 12% of that from Wind for all stages. Under uncertainty in R&D, the median energy production is the same as the deterministic solution. But the 90% range runs from zero on the low end (after Stage 3) to as high as 70% that of Wind. The difference between the deterministic and stochastic results for

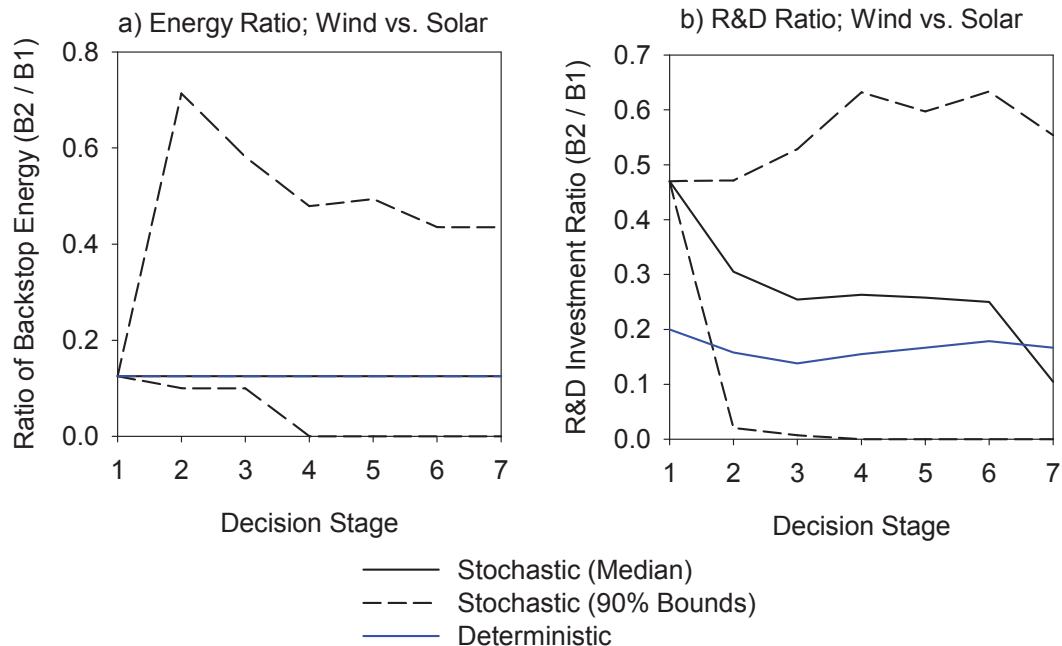


Figure 8: Optimal ratios of backstop energy production and R&D investment (Solar / Wind). For stochastic case, median and 90% bounds are shown in black, and deterministic solution is shown in blue.

the R&D ratio is more striking (Fig. 7b). The optimal R&D for Solar in the absence of uncertainty is roughly 20% of that for Wind. Under uncertainty, the optimal first stage R&D ratio is significantly higher, nearly 50%. The median R&D ratio remains well above the deterministic solution for all except the last stage. The 90% range for the R&D ratio is quite broad, from zero to 60%.

Finally, we use the results of simulations of size $N=1000$ sample paths (after convergence) of random R&D shocks to examine the question of how to respond to a below average result from R&D into Solar. Because the random returns to R&D and the decisions are all continuous, we use illustrative definitions for a “poor outcome” from R&D and for “eventual success” of a technology. Different threshold values could be used, but the qualitative results remain the same. Here, we assume that a random draw for the multiplicative shock to the returns to R&D into Solar of less than 0.5 (where 1.0 defines the mean outcome) defines a “poor outcome”. We also define “eventual success” of Solar as a sample path for which the backstop energy produced from Solar in Stage 7 is greater than 1 (the median outcome over all 1000 samples is 0.98). This set of sample paths will consist of the samples for which there is more substitution of Solar for Wind than occurs in the deterministic solution.

Using these thresholds to operationalize the concepts, within a sample set of 1000, there were 387 for which the random shock to R&D for Solar was less than 0.5 in the first stage, or a probability of 0.387. We wish to consider the optimal R&D into Solar in Stage 2, conditional on observing this “poor outcome” from Stage 1, and also the conditional probability of eventual success given this outcome from the first stage. The model gives a conditional optimal R&D decision for Solar in Stage 2 of \$634M, not significantly less than the unconditional optimal decision of \$710M. One poor draw is not sufficient to justify major reductions in R&D into

Solar. Conditional on a poor outcome in Stage 1 as defined here, the probability of eventual success of Solar is 0.018. This is considerably less than the unconditional probability of success across all 1000 samples, 0.071, but may still be large enough to justify continued R&D for another stage.

Of course, repeated poor outcomes over many stages should not lead to continued R&D investment at the same level. In this model, if the outcomes of the first three stages are poor outcomes (less than half the expected value), then the optimal R&D investment into Solar from Stage 4 onwards is effectively zero.

5. Concluding Discussion

This paper has addressed the question of optimal R&D portfolios into low carbon energy technologies in the context of climate change. We developed both simple analytical and detailed numerical models to represent two substitutable low-carbon energy technologies, each of which can have costs reduced through R&D. Using both approaches, we have demonstrated that given two technologies with different initial costs and uncertainty in the outcomes of R&D into the higher cost technology, the optimal R&D investment into the higher cost technology is increasing in both the variance and the positive skewness of the R&D uncertainty.

In contrast with other approaches that model specific technologies in more detail to examine the optimal portfolio question, we have deliberately focused on a stylized model of two technologies and performed sensitivity analyses to explore the general conditions that justify R&D investment into a higher cost substitute. Nevertheless, we do present one analysis in which parameters are calibrated to be consistent with observations of wind and solar patent citations. Using that example, we illustrate that significant R&D investment into solar is

justified – the optimal portfolio would divide the R&D budget into roughly 2/3 to wind and 1/3 to solar in the first stage. We also illustrate in that example that a poor outcome from R&D into solar in the first stage does not lead to a large reduction in the optimal solar R&D in stage 2. Repeated poor outcomes over three or more stages must be observed before abandoning solar is, however, justified.

It is important to stress the related questions that we explicitly do not address, because the model used here is not appropriate for them. One key question is the socially optimal division of R&D between the public and private sectors. This question requires more detailed modeling of spillovers and knowledge externalities than captured in ENTICE. Nor do we address how R&D into energy technologies should be, or will be, divided among different nations. Clearly, energy technologies exist in a global market, and technological advances achieved in one country reduce the costs of that technology to consumers in other countries. However, this issue involves complex issues of economic competitiveness. The results have focused only on the socially optimal total global R&D, without consideration of public/private or international divisions of effort. Finally, we have examined here the case of two easily substitutable energy technologies. Some technologies are complementary in order to make the energy system function reliably, and the implications for R&D portfolios over a more general set of technologies is left for future work.

Our main contribution to the debate over energy technology R&D is conceptual. Given two substitutable technologies, where one is a higher cost today, greater uncertainty in future R&D outcomes justifies more R&D into that technology. In particular, if one believes that the distribution of R&D outcomes is highly positively skewed, then considerable R&D may be justified. A direct consequence of a positively skewed distribution is that the vast majority of

outcomes from individual R&D projects will be well below average. Thus a poor outcome from R&D may not necessarily be reason to abandon future efforts; rather many failures may be expected along the road to long-shot technology breakthroughs. This perspective lends support to several recent efforts in the U.S. Department of Energy, including ARPA-E and the Sunshot Program. An efficient R&D portfolio will therefore allocate R&D both to currently competitive technologies and also to longshots, in proportion to the expected benefits if successful after all.

Acknowledgements

The authors gratefully acknowledge the financial support of U.S. National Science Foundation, grant # 0825915.

References

- Acemoglu, D., 2002. Directed technical change. *Review of Economic Studies* 69 (2002) 781 – 809.
- Acemoglu, D., 1998. Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality. *The Quarterly Journal of Economics* (1998) 113 (4): 1055-1089.
- Baker, E., Adu-Bonnah, K., 2008. Investment in risky R&D programs in the face of climate uncertainty. *Energy Economics* 30 (2008) 465 – 486.
- Baker, E., Solak, S., 2011. Climate change and optimal energy technology R&D policy. *European Journal of Operational Research* 213 (2011) 442 – 454.
- Baker, E., Chon, H., Keisler, J. M., 2008. Advanced Nuclear Power: Combining Economic Analysis with Expert Elicitations to Inform Climate Policy. (August 8, 2008). Available at SSRN: <http://ssrn.com/abstract=1407048> or <http://dx.doi.org/10.2139/ssrn.1407048> Working Paper.
- Baker, E., Chon, H., Keisler, J., 2009a. Advanced solar R&D: applying expert elicitations to inform climate policy. *Energy Economics* 31, S37–S49.
- Baker, E., Chon, H., Keisler, J., 2009b. Carbon capture and storage: combining expert elicitations to inform climate policy. *Climatic Change* 96 (3), 379 – 408.
- Baker, E., Clarke, L., and Weyant, J., 2006. Optimal Technology R&D in the Face of Climate Uncertainty. *Climatic Change* 75: 157 – 180.
- Barbose, G. L., Weaver, S., Darghouth, N.R., 2014. Tracking the Sun VII: An Historical Summary of the Installed Price of Photovoltaics in the United States from 1998-2013. Lawrence Berkeley National Laboratory. Report LBNL-6858E. Available at: <http://emp.lbl.gov/sites/all/files/lbnl-6858e.pdf>.
- Bellman, R., 2003. *Dynamic programming*. Dover, New York.
- Bertsekas, D., Tsitsiklis, J., 1996. *Neuro dynamic programming*. Athena Scientific, Belmont, MA.
- Bosetti, V., Tavoni, M., 2009. Uncertain R&D, backstop technology and GHGs Stabilization. *Energy Economics* 31 (2009) S18 – S26.
- Bosetti, V., Carraro, C., Galeotti, M., Massetti, E., Tavoni, M., 2006. WITCH: A world induced technical change hybrid model. *The Energy Journal*. Hybrid Modeling of Energy-Environment Policies: Reconciling Bottom-up and Top-down.

- Bosetti, V., Carraro, C., Massetti, E., Sgobbi, A., Tavoni, M., 2009. Optimal energy investment and R&D strategies to stabilize atmospheric greenhouse gas concentrations. *Resource and Energy Economics* 31 (2009) 123 – 137.
- Brenkert, A.L., Kim, S.H., Smith, A.J., Pitcher, H.M., 2003. Model Documentation for the MiniCAM. Report PNNL-14337, Pacific Northwest National Laboratory.
- Buonanno, P., Carraro, C., Galeotti, M., 2003. Endogenous induced technical change and the costs of Kyoto. *Resource and Energy Economics* 25 (2003) 11 – 34.
- Church, S., 2011. Evergreen Solar Seeks Bankruptcy With Plans to Sell Itself. Bloomberg News. August 15, 2011. Available at: <http://www.bloomberg.com/news/articles/2011-08-15/evergreen-solar-seeks-bankruptcy-protection-with-debt-of-486-5-million>.
- Energy Information Administration, 2015. Levelized Cost and Levelized Avoided Cost of New Generation Resources in the Annual Energy Outlook 2015. Available at: http://www.eia.gov/forecasts/aeo/electricity_generation.cfm.
- Fischer et al., 2013; “Environmental and Technology Policy Options in the Electricity Sector: Interactions and Outcomes” RFF Discussion Paper 13-20, December 2013 <http://www.rff.org/Publications/Pages/PublicationDetails.aspx?PublicationID=22362>.
- Fischer, C., Newell, R.G., 2008. Environmental and technology policies for climate mitigation. *Journal of Environmental Economics and Management* 55 (2008) 142 – 162.
- Gillingham, K., Newell, R.G., Pizer, W.A., 2008. Modeling endogenous technological change for climate policy analysis. *Energy Economics* 30 (2008) 2734 – 2753.
- Goeschl, T., Perino, G., 2009. On backstops and boomerangs: Environmental R&D under technological uncertainty. *Energy Economics* 31 (2009) 800 – 809.
- Goulder, L.H., Mathai, K., 2000. Optimal CO₂ abatement in the presence of induced technological change. *Journal of Environmental Economics and Management* 39, 1–38.
- Goulder, L.H., Schneider, S.H., 1999. Induced technological change and the attractiveness of CO₂ abatement policies. *Resource and Energy Economics* 21, 211–253.
- Green, M.A., 2000. Photovoltaics: technology overview. *Energy Policy* 28 (2000) 989 – 998.
- Gritsevski, A., Nakicenovic, N., 2000. Modeling uncertainty of induced technological change. *Energy Policy* 28 (2000) 907 – 921.
- Hung, N.M., Quyen, N.V., 1993. On R&D timing under uncertainty: the case of exhaustible resource substitution. *Journal of Economic Dynamics and Control* 17 (1993) 971 – 991.

- Jaffe, A.B. and Trajtenberg, M. 2002. *Patents, Citations, & Innovations: A Window on the Knowledge Economy*, MIT Press, Cambridge, MA
- Kennedy, C., 1964. Induced Bias in Innovation and the Theory of Distribution. *The Economic Journal* 74 (295) 541-547.
- Kamien, M.I., Schwartz, N.L., 1968. Optimal “induced” technical change. *Econometrica*. 36 (1) 1 – 17.
- Lanjouw, J.O., and Schankerman, M. 2004. Patent Quality and Research Productivity: Measuring Innovation with Multiple Indicators. *Economic Journal* 114 (495): 441-465.
- Löschel, A., 2002. Technological change in economic models of environmental policy: a survey. *Ecological Economics* 43 (2–3) 105-126.
- Messner, S., Golodnikov, A., Gritsevskiy, A., 1996. A stochastic version of the dynamic linear programming model MESSAGE III. *Energy* 21 (9).
- Nordhaus, W., Boyer, J., 2000. *Warming the world: economic modeling of global warming*. MIT Press, Cambridge, MA
- Pakes, Ariel, (1986), “Patents as options: Some estimates of the value of holding European patent stocks,” *Econometrica* 54 (4) 755-84.
- Pizer, W.A., Popp, D., 2008. Endogenizing technological change: Matching empirical evidence to modeling needs. *Energy Economics* 30 (2008) 2754 – 2770.
- Popp, D.C., 2001. The effect of new technology on energy consumption. *Resource and Energy Economics* 23 (2001) 215 – 239.
- Popp, D. 2002. Induced Innovation and Energy Prices. *American Economic Review* 92 (1): 160-180.
- Popp, D., 2004. ENTICE: endogenous technological change in the DICE model of global warming. *Journal of Environmental Economics and Management* 48 (2004) 742 – 768.
- Popp, D., 2006a. ENTICE-BR: The effects of backstop technology R&D on climate policy models. *Energy Economics* 28 (2006) 188 – 222.
- Popp, D., 2006b. Comparison of climate policies in the ENTICE-BR model. *The Energy Journal*. Endogenous Technological Change and the Economics of Atmospheric Stabilization Special Issue.
- Popp, D., Newell, R.G., and A.B. Jaffe (2010). “Energy, the environment, and technological change,” *Handbook of the Economics of Innovation*: vol. 2, Bronwyn Hall and Nathan Rosenberg, eds., Academic Press/Elsevier, 2010, 873-937.

- Popp, D., Santen, N., Fisher-Vanden, K., Webster, M., 2013. Technology Variation vs. R&D Uncertainty: What Matters Most for Energy Patent Success? *Resource and Energy Economics* 35 (4): 505-533.
- Powell, W., 2007. *Approximate dynamic programming: solving the curses of dimensionality*. Wiley-Blackwell, New York.
- Pugh, G., Clarke, L., Marlay, R., Kyle, P., Wise, M., McJeon, H., Chan, G., 2011. Energy R&D portfolio analysis based on climate change mitigation. *Energy Economics* 33 (2011) 634 – 643.
- Romer, P., 1990. Endogenous technological change. *Journal of Political Economy* 94, 1002 – 1037.
- Schankerman, Mark (1998), “How valuable is patent protection? Estimates by technology field,” *RAND Journal of Economics* 29 (1) 77-107.
- Schankerman, Mark and Ariel Pakes (1986), “Estimates of the value of patent rights in European countries during the post-1990 period,” *Economic Journal* 96 (384) 1052-1076.
- Sue Wing, I., 2006. Representing induced technological change in models for climate policy analysis. *Energy Economics* 28 (5–6) 539-562.
- Sunshot, 2012. Sunshot Vision Study. February, 2012. U.S. Department of Energy. Available at: <http://energy.gov/eere/sunshot/sunshot-vision-study>.
- Tsur, Y., Zemel, A., 2005. Scarcity, growth, and R&D. *Journal of Environmental Economics and Management* 49 (2005) 484 – 499.
- U.S. Department of Energy, 2015. *Quadrennial Energy Review: Energy Transmission, Storage, and Distribution Infrastructure*.
http://energy.gov/sites/prod/files/2015/05/f22/QER%20Full%20Report_0.pdf
- U.S. News & World Report, 2012. Should the Government Invest in Green Energy? U.S. News & World Report. January 18, 2012. Available at: <http://www.usnews.com/debate-club/should-the-government-invest-in-green-energy>.
- Webster, M.D., Santen, N.R., Parpas, P., 2012. An approximate dynamic programming framework for modeling global climate policy under decision-dependent uncertainty. *Computational Management Science* 9: 339–362.
- Weiner, R., 2012. Solyndra Explained. The Washington Post. June 1, 2012. Available at: www.washingtonpost.com/blogs/the-fix/post/solyndra--explained/2012/06/01/gJQAig2g6U_blog.html

Appendix A: Model Documentation

Original ENTICE Model Equations

This section documents the parameters, variables, and equations of the original ENTICE-BR model (Popp 2006a; 2006b).

Exogenous variables and parameters

t	time
L_t	population at time t , also equal to labor inputs
L_0	initial population level
$g_{L,t}$	growth rate of population
$g_{L,0}$	initial value of the growth rate of the population
d_L	rate of decline of $g_{L,t}$
D_t	pure rate of time preference discount factor
r_0	initial value of the pure rate of time preference
g_r	growth rate of social time preference
A_t	total factor productivity
A_0	initial value of total factor productivity
$g_{A,t}$	growth rate of total factor productivity
$g_{A,0}$	initial value of the growth rate for TFP
d_A	rate of decline for $g_{A,t}$
γ	elasticity of output with respect to capital
β	elasticity of output with respect to energy/carbon inputs
Φ_t	ratio of carbon emissions per unit of carbon services
g_t^z	growth rate of Φ_t per decade
δ^z	rate of decline of g_t^z
$\zeta_1, \zeta_2, \zeta_3$	parameters of the long-run carbon supply curve
markup	energy services price markup
CumC*	total carbon resources available
δ	rate of depreciation of physical capital stock
δ_H	rate of depreciation of energy knowledge stock
crowdout	percentage of overall R&D crowded out by energy R&D
a, b, ϕ	parameters of the innovation possibilities frontier
η	effect of backstop energy knowledge on backstop price
α_H	scaling factor for the stock of energy knowledge
α_Φ	percentage of exogenous carbon intensity reduction
ρ_H	substitution parameter between energy and knowledge
ρ_B	substitution parameter between fossil fuels and backstop energy
LU_t	land-use carbon emissions
LU_0	initial land-use carbon emissions
δ_{LU}	rate of decline of land-use carbon emissions
$\phi_{11}, \phi_{12}, \phi_{21}$	parameters of the carbon transition matrix
$\phi_{22}, \phi_{23}, \phi_{32}, \phi_{33}$	
O_t	increase in radiative forcing from anthropogenic non-carbon sources
$\sigma_1, \sigma_2, \sigma_3$	temperature dynamics parameters
θ_1, θ_2	parameters of the damage function
$4.1/\lambda$	climate sensitivity

Endogenous variables

U_t	utility in period t
c_t	per capita consumption
Q_t	output (trillions of US dollars)
Ω_t	damages from climate change
μ_t	emissions control rate in DICE model
K_t	physical capital stock (trillions of US \$)
E_t	energy inputs
$p_{F,t}$	price of fossil fuels
$p_{B,t}$	price of backstop energy
F_t	fossil fuel/carbon inputs, also equal to CO ₂ emissions
B_t	backstop energy, in carbon ton equivalents (CTE)
q_F	marginal cost of fossil fuel extraction
$CumC_t$	cumulative carbon extractions by year t
I_t	investment in physical capital
C_t	total consumption
$H_{E,t}$	stock of energy efficiency knowledge
$H_{B,t}$	stock of backstop energy knowledge
$R_{E,t}$	energy R&D
EM_t	carbon emissions
$M_{A,t}$	atmospheric CO ₂ concentration
$M_{U,t}$	upper oceans/biosphere CO ₂ concentration
$M_{L,t}$	lower oceans CO ₂ concentration
$FORCE_t$	radiative forcing, increase over preindustrial level
T_t	atmospheric temperature, increase over 1900 level
TL_t	lower ocean temperature, increase over 1900 level

Model Equations

The ENTICE model maximizes per capita utility, defined in Eq. (A1) below, subject to a set of environmental and economic constraints. Economic constraints are represented by Eqs. (A2)-(A18), and the environmental constraints are represented by Eqs. (A19)-(A28).

$$\max V = \sum_{t=0}^T U[c_t, L_t] D_t \quad (\text{A1})$$

Economic Constraints

$$U_t = L_t \log(C_t/L_t) \quad (\text{A2})$$

$$D_t = \prod_{\tau=0}^t [1 + r_0 e^{-g_r \tau}]^{10} \quad (\text{A3})$$

$$Q_t = \Omega_t(A_t K_t^\gamma L_t^{1-\gamma-\beta} E_t^\beta) - p_{F,t} F_t - p_{B,t} B_t \quad (\text{A4})$$

$$K_t = I_t - 4 * \text{crowdout} * (R_{E,t} + R_{B,t}) + (1 - \delta) K_{t-1} \quad (\text{A5})$$

$$L_t = L_0 e^{g_{L,t}} \quad (\text{A6})$$

$$g_{L,t} = (g_{L,0}/d_L) * (1 - e^{-d_L t}) \quad (\text{A7})$$

$$A_t = A_0 e^{g_{A,t}} \quad (\text{A8})$$

$$g_{A,t} = (g_{A,0}/d_A) * (1 - e^{-d_A t}) \quad (\text{A9})$$

$$E_t = \left[\alpha_H H_{E,t}^{\rho_H} + \left(\left(\frac{F_t}{\alpha_\Phi \Phi_t} \right)^{\rho_B} + B_t^{\rho_B} \right)^{\rho_H/\rho_B} \right]^{1/\rho_H} \quad \rho < 1 \quad (\text{A10})$$

$$\Phi_t = \exp \left[\left(\frac{g_t^z}{\delta^z} \right) (1 - \exp(-\delta^z t)) \right] \quad (\text{A11})$$

$$P_F = q_F + \text{markup} \quad (\text{A12})$$

$$q_F = \zeta_1 + \zeta_2 [\text{Cum}C_t / \text{Cum}C^*]^{\zeta_3} \quad (\text{A13})$$

$$\text{Cum}C_t = \text{Cum}C_{t-1} + 10 * F_t \quad (\text{A14})$$

$$F_t < 0.1 * (\text{Carbmax} - \text{Cum}C_t) / 10 \quad (\text{A15})$$

$$p_{B,t} = \frac{p_{B,0}}{H_{B,t}^\eta} \quad (\text{A16})$$

$$H_{i,t} = h(R_{i,t}) + (1 - \delta_H)H_{i,t-1}, \quad i = E, B \quad (\text{A17})$$

$$h(R_{i,t}) = aR_{i,t}^b H_{i,t}^\Phi, \quad i = E, B \quad (\text{A18})$$

$$Q_t = C_t + I_t + R_{E,t} + R_{B,t} \quad (\text{A19})$$

Environmental Constraints

$$LU_t = LU_0(1 - \delta_{LU})^t \quad (\text{A20})$$

$$EM_t = F_t + LU_t \quad (\text{A21})$$

$$M_{A,t} = 10 * EM_t + \phi_{33}M_{L,t-1} + \phi_{23}M_{U,t-1} \quad (\text{A22})$$

$$M_{L,t} = \phi_{11}M_{A,t-1} + \phi_{21}M_{U,t-1} \quad (\text{A23})$$

$$M_{U,t} = \phi_{12}M_{A,t-1} + \phi_{22}M_{U,t-1} + \phi_{32}M_{L,t-1} \quad (\text{A24})$$

$$FORCE_t = 4.1 * \log(M_{A,t}/596.4)/\log(2) + O_t \quad (\text{A25})$$

$$O_t = -0.1965 + 0.13465t, \quad t < 11, \quad O_t = 1.15, \quad t \geq 11 \quad (\text{A26})$$

$$T_t = T_{t-1} + \sigma_1 FORCE_t - \lambda T_{t-1} - \sigma_2(T_{t-1} - TL_{t-1}) \quad (\text{A27})$$

$$TL_t = TL_{t-1} + \sigma_3(T_{t-1} - TL_{t-1}) \quad (\text{A28})$$

$$\Omega_t = 1/(1 + a_1 T_t + a_2 T_t^2) \quad (\text{A29})$$

Modifications to ENTICE Model

The modifications to the ENTICE-BR model introduce two substitutable backstop technologies instead of one, and introduce two knowledge stocks for each backstop, one for creating new knowledge and one for cost reduction.

Exogenous parameters

ρ_κ	substitution parameter between the two backstops
$p_{B1,0}$	initial price of backstop 1
$p_{B2,0}$	initial price of backstop 2

Endogenous variables

$B_{1,t}$	energy from Backstop 1 (CTE)
$B_{2,t}$	energy from Backstop 2 (CTE)
$p_{B1,t}$	price of Backstop 1 at time t
$p_{B2,t}$	price of Backstop 2 at time t
$H_{B1,t}^K$	knowledge stock for Backstop 1 for new knowledge production
$H_{B2,t}^K$	knowledge stock for Backstop 2 for new knowledge production
$H_{B1,t}^C$	knowledge stock for Backstop 1 for cost reduction
$H_{B2,t}^C$	knowledge stock for Backstop 2 for cost reduction

Changes to model equations

Include both backstops in the energy production function by replacing Eq. (A10) with:

$$E_t = \left[\alpha_H H_{E,t}^{\rho_H} + \left(\left(\frac{F_t}{\alpha_\Phi \Phi_t} \right)^{\rho_B} + (B_{1,t}^{\rho_\kappa} + B_{2,t}^{\rho_\kappa})^{\rho_B/\rho_\kappa} \right)^{\rho_H/\rho_B} \right]^{1/\rho_H}. \quad (\text{A30})$$

The production function for the economy must include the costs of using both backstops; replace Eq. (A4) with

$$Q_t = \Omega_t (A_t K_t^\gamma L_t^{1-\gamma-\beta} E_t^\beta) - p_{F,t} F_t - p_{B1,t} B_{1,t} - p_{B2,t} B_{2,t}. \quad (\text{A31})$$

In the knowledge stock dynamics equations, knowledge stocks for creating new knowledge depreciate, but stocks for cost reduction do not depreciate. Replace Eq. (A17) with the following two equations:

$$H_{B,t}^K = h(R_{B,t}) + (1 - \delta_B) H_{B,t-1}^K, \quad B = B_1, B_2 \quad (\text{A32})$$

$$H_{B,t}^C = h(R_{B,t}) + H_{B,t-1}^C, \quad B = B_1, B_2 \quad (\text{A33})$$

In the innovations possibility frontier, the knowledge stock for new knowledge is used; replace Eq. (A18) with:

$$h(R_{B,t}) = aR_{B,t}^b(H_{B,t}^K)^{\Phi_B}, \quad B = B_1, B_2. \quad (\text{A34})$$

In the backstop cost function, the knowledge stock for cost reduction is used; replace Eq. (A16) with:

$$p_{B,t} = \frac{p_{B,0}}{(H_{B,t}^C)^\eta} \quad (\text{A35})$$

Parameter value assumptions

In Table A.1, we give the values of all constants in the above equations.

Approximate Dynamic Programming Implementation

Here, we briefly present the solution algorithm for the Approximate Dynamic Programming (ADP) version of ENTICE-BR. More details can be found in Webster *et al.* (2012).

Our approximate dynamic programming (ADP) algorithm for solving this problem is shown in Algorithm 1. ADP is a family of methods (e.g., Bertsekas and Tsitsiklis, 1996; Powell, 2007) that approximates the value function in each stage by adaptively sampling the state space to focus on higher expected value states until the value function converges. One critical advantage of forward sampling is that this enables a straightforward representation of decision-dependency. Two critical design choices in any efficient ADP algorithm are 1) the sampling strategy, and 2) the value function approximation.

Our solution algorithm consists of two phases. In phase I, the bootstrap phase, we use Latin Hypercube Sampling to explore both the action space over all stages and the R&D result shock space. These sample paths are simulated forward, and the resulting Bellman values for the sample states and actions are saved for each decision stage. The full set of these samples of the value function are used to produce the first estimate of the value function approximation for each decision stage, using either of the two methods described below.

In phase II, we randomly sample the cost shock in each period to obtain a sample path, and choose the optimal action in each stage using the current value function approximations for the value of the next state, and the simulated ENTICE-BR equations to obtain the current reward. The overall sampling approach is an efficient (stratified) pure explore strategy in Phase I and a pure exploit strategy in Phase II.

We employ a parametric approach, using an iterative least squares regression method, and approximating the value function as a linear function of the prices

Algorithm 1: DICE Approximate Dynamic Programming Algorithm

Input: Decision stages N , bootstrap iterations bs , possible controls μ , uncertainty variable $\theta \sim N(1, \sigma)$, system state $s_0 \in S$ at time t_0 , system state transition equations $F(\mu, \theta)$, convergence criterion, $\bar{\epsilon}$

Phase I Initialization-Bootstrap: While $i \leq bs$,

1. Forward Pass

Loop over t from **1** to N , Latin Hypercube Sampling from θ and μ and set current reward as:

$$R_t(s_i) = U(c_t, L_t)(1 + \rho_t)^{-1}.$$

2. Backward Pass

Loop over t from N to **1**, setting the Bellman Value as:

$$v_t(s_i) = (R_t(s_i) + v_{t+1}(y_i|s_i))$$

where y_i is the sampled next system state resulting from μ_t and θ_t , and v_N is a pre-defined terminal value.

3. Construct First Estimate of Value Function: When $i = bs$, use OLS to set:

$$\hat{v}_i(s) = \Phi(s)r_0,$$

where Φ is a row vector of basis functions and r_0 is a column vector of coefficients that solves:

$$\min_{r_0} \sum_{s_i} (\hat{v}_i(s_i) - \Phi(s_i)r_0)^2.$$

for all sample states s_i .

Phase II Main Loop-Optimization: While $i > bs$,

1. Forward Pass

Loop over t from 1 to N , sampling θ randomly and sampling controls μ that achieve:

$$\max_{\mu} [R_t(s_i) + E \{v_{t+1}(y_i|s_i)\}]$$

where

$$E \{v_{t+1}(y_i|s_i)\} = \hat{v}_{t+1}(\mu_t, \theta_t).$$

Set current reward, $R_t(s_i)$, as in Phase I.

2. Backward Pass

Loop over t from N to **1**, setting the new Bellman Value as:

$$v_t(s_i) = (R_t(s_i) + \hat{v}_{t+1}(y_i|s_i))\delta$$

where y_i is the sampled next system state.

Update r_i using a Bellman Error routine:

$$r_{i+1} = r_i - \gamma_i \epsilon_i \nabla r_i$$

where γ_i is a predefined smoothing parameter and

of the two backstops $(p_{B1,t}, p_{B2,t})$. The approximation of the value function is

$$\hat{v}_t(s) = \Phi(s)r \tag{A35}$$

where Φ is a row vector of basis functions and r is a column vector of coefficients that solves,

$$\min_r \sum_{s_i} (\hat{v}_t(s_i) - \Phi(s_i)r)^2.$$

for all sample states s_i . Given an initial estimate of the coefficient vector r from the bootstrap phase, we iteratively improve the estimate using a Bellman error approach.

Sensitivity of Results to Number of Decision Stages

One key assumption in the dynamic programming formulation is the number of decision stages. The deterministic version of ENTICE-BR solves for 35 decadal time-steps. Even with the computational efficiency of the ADP implementation, convergence of a 35-stage model would be prohibitive to solve for the large number of scenarios explored in the paper. We therefore reduce the DP formulation to a smaller number of decision stages, each of which consists of several decades. This raises the question of how the results vary for differing numbers of stages.

Here we briefly present a sample of results for several different numbers of stages. Figure 1 shows the ratio of optimal stage 1 R&D for three initial ($t = 0$) backstop prices of \$1300, \$1500, and \$1800. Each line shows the ratios as a function of initial backstop price for different numbers of decisions stages ($N = 3, 5, 7, 9, 11$). There are two main features to note. First, relatively more R&D investment into the risky technology (Backstop 2) is optimal for shorter decision periods (i.e., larger numbers of decision stages). This is because with more future stages to reinvest in case of an early success, the expected marginal benefit of investment increases. However, it is also true that the qualitative results of the paper, such as the sensitivity to the initial backstop price, is robust to the number of stages. Notice that the curves all follow the same trend, and are simply vertically shifted from one another. Given the robustness of this and other key results, we have chosen $N = 7$ stages as a compromise between representing the multi-stage characteristics of this problem, and the computational burden of achieving convergence for every case explored in the paper.

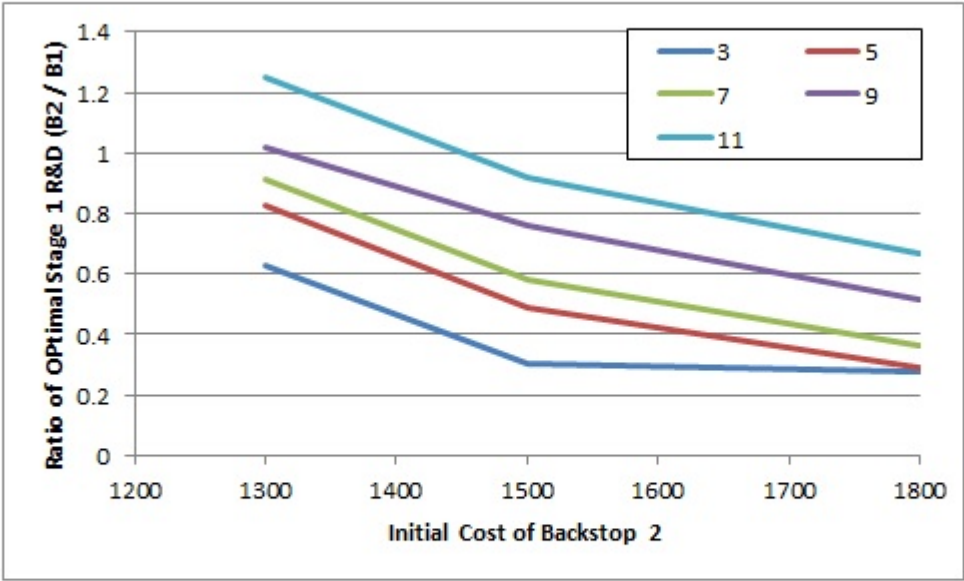


Figure 1: Sensitivity of Optimal R&D Ratio to Number of Decision Stages

Table 1: ENTICE-BR Parameter Assumptions

Parameter	Value	Description
r_0	0.03	Initial rate of social time preference
$g_{L,0}$	0.157	Initial population growth rate
d_L	0.222	Decline rate of population growth
L_0	5632.7	1990 global population
γ	0.30	Elasticity of output with respect to capital
$g_{A,0}$	0.038	Initial growth rate for productivity
d_A	0.000001	Decline rate of productivity growth
$A(0)$	0.01685	Initial level of total factor productivity
θ_1	-0.0045	Damage coefficient linear term
θ_2	0.0035	Damage coefficient quadratic term
δ_K	0.1	Depreciation rate on capital stock
K_0	51.3	Initial year value of capital stock
β	0.0816	elasticity of output with respect to energy/carbon inputs
Φ_0	-0.1549	Initial value of Φ_t
δ^z	0.2396	rate of decline of g_t^z
ζ_1	0	constant term for long-run carbon supply curve
ζ_2	976.29	coefficient for long-run carbon supply curve
ζ_3	4	exponent for long-run carbon supply curve
CumC*	6000	maximum cumulative fossil fuel resource
δ_H	0.01	rate of depreciation of energy knowledge stock
crowdout	1.5	fraction of overall R&D crowded out by energy R&D
a_E	0.0264	scaling parameter of the IPF for energy efficiency
b_E	0.20	returns to R&D parameter of the IPF for energy efficiency
ϕ_E	0.54	return to knowledge stock parameter of the IPF for e.e.
a_B	0.0170	scaling parameter of the IPF for backstop
b_B	0.10	returns to R&D parameter of the IPF for backstop
ϕ_B	0.54	return to knowledge stock parameter of the IPF for backstop
η	0.4	effect of backstop energy knowledge on backstop price
ρ_H	0.38	substitution parameter between energy and knowledge
ρ_B	0.5423	substitution parameter between fossil fuels and backstop energy
ρ_κ	0.80	substitution parameter between backstops
δ_{LU}	0.1	Decline rate of land-use emissions
$LU(0)$	1.128	Initial year land-use emissions
ϕ_{11}	0.66616	Carbon-cycle transition matrix
ϕ_{12}	0.33384	Carbon-cycle transition matrix
ϕ_{21}	0.27607	Carbon-cycle transition matrix
ϕ_{22}	0.60897	Carbon-cycle transition matrix
ϕ_{23}	0.11496	Carbon-cycle transition matrix
ϕ_{32}	0.00422	Carbon-cycle transition matrix
ϕ_{33}	0.99578	Carbon-cycle transition matrix
$M_{A,0}$	735	Carbon concentration in atmosphere 1990
$M_{U,0}$	781	Carbon concentration in upper ocean 1990
$M_{L,0}$	19230	Carbon concentration in deep ocean 1990
σ_1	0.226	Climate equation coefficient for surface layer
CS	2.9078	Climate sensitivity
σ_2	0.440	Transfer coefficient for surface to deep layer
σ_3	0.02	Transfer coefficient for deep to surface layer
T_0	0.43	1985 surface temperature change from 1900
TL_0	0.06	1985 deep ocean temperature change from 1900