Energy Technology Innovation Policy

Electricity Technology Investments Under Solar RD&D Uncertainty:

How Interim Learning and Adaptation Affect the Optimal Decision Strategy

Nidhi R. Santen and Laura Diaz Anadon

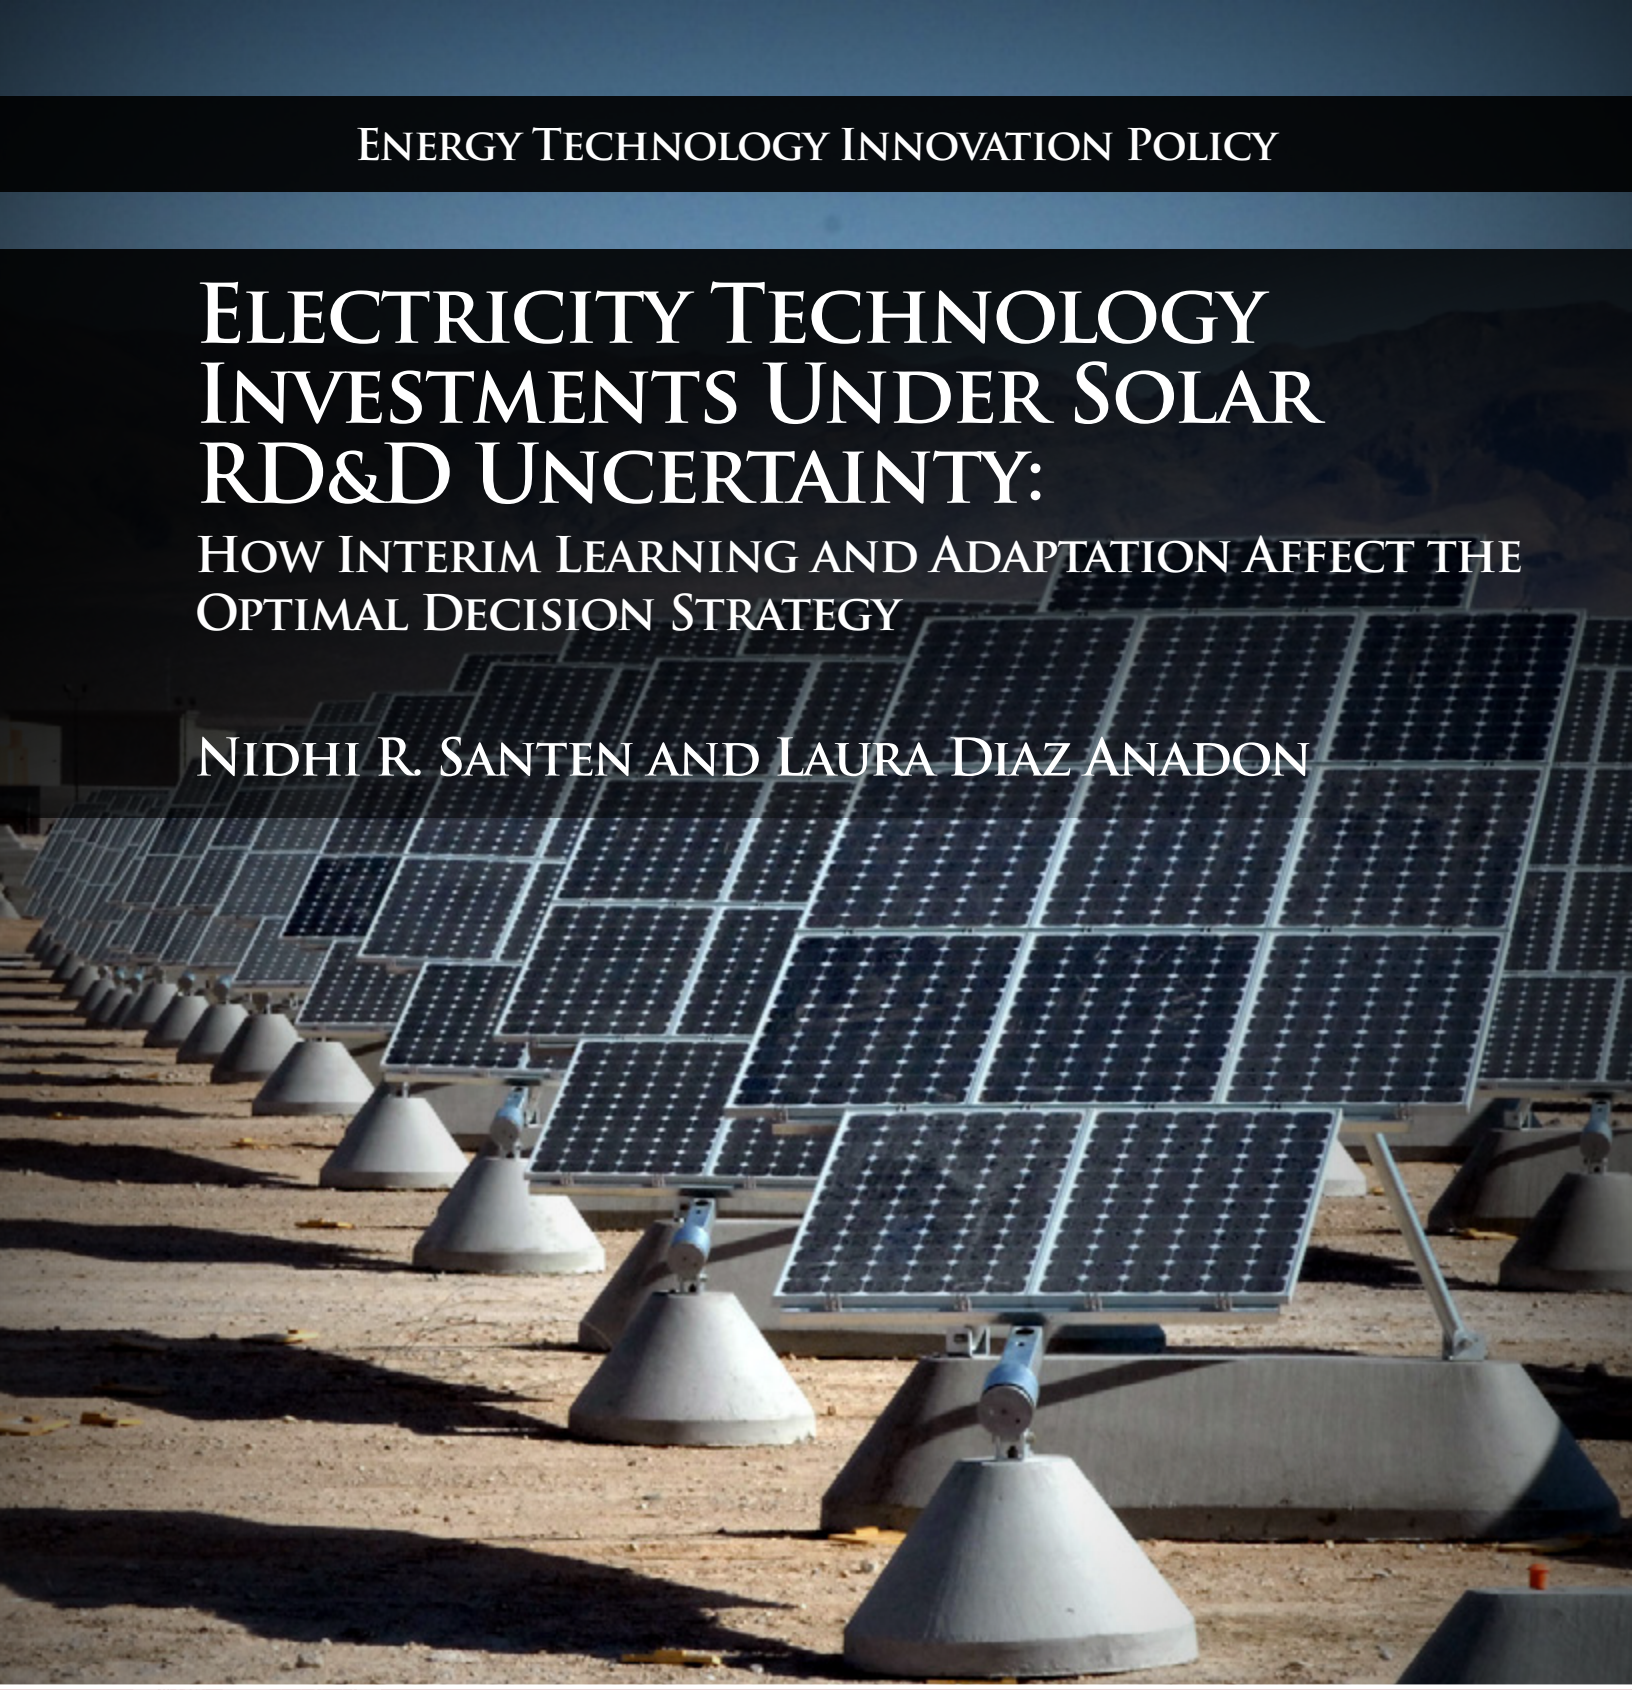

HARVARD Kennedy School
Belfer Center for Science and International Affairs

December 2014

Belfer Center for Science and International Affairs
Energy Technology Innovation Policy research group, a joint effort between the Science, Technology, and Public Policy Program and the Environment and Natural Resources Program.

Discussion Paper #2014-10

John F. Kennedy School of Government
Harvard University
79 JFK Street
Cambridge, MA 02138

Nidhi R. Santen[1] and **Laura Diaz Anadon**[2]

[1] Corresponding Author. Postdoctoral Research Fellow, Energy Technology Innovation Policy Research Group, Belfer Center for Science and International Affairs, John F. Kennedy School of Government, Harvard University, 79 John F. Kennedy Street, Mailbox 53, Cambridge, MA 02138, e-mail: nidhi_santen@hks.harvard.edu, ph: 617-495-1853, fax: 617-495-8963

[2] Assistant Professor of Public Policy; Co-Principal Investigator, Energy Technology Innovation Policy Research Group; Associate Director, Science, Technology, and Public Policy Program, Belfer Center for Science and International Affairs, John F. Kennedy School of Government, Harvard University, 79 John F. Kennedy Street, Box 53, Cambridge, MA 02138

Santen, Nidhi R. and Laura Diaz Anadon. "Electricity technology investments under solar RD&D uncertainty: How interim learning and adaptation affect the optimal decision strategy." Discussion Paper 2014-10, Belfer Center for Science and Inernational Affairs, Harvard Kennedy School. December 2014.

Cover image: The largest photovoltaic solar power plant in the United States is becoming a reality at Nellis Air Force Base. When completed in December, the solar arrays will produce 15 megawatts of powe. (USAF photo)

Electricity Technology Investments Under Solar RD&D Uncertainty:

How interim learning and adaptation affect the optimal decision strategy

Nidhi R. Santen and Laura Diaz Anadon

HARVARD Kennedy School
Belfer Center for Science and International Affairs

December 2014

Abstract

We present a new modeling framework for studying optimal generating capacity and public RD&D investments in the electricity sector under decision-dependent RD&D uncertainty and learning. We show that when uncertainty and learning features are omitted, as is typically done in practice, the investment strategy can be considerably different from the optimal solution. We use a bottom-up stochastic power generation capacity expansion model with uncertain endogenous RD&D-based technical change, and focus on solar PV RD&D for its current prominent role in the U.S. national energy and climate policy discussion. Uncertainty in the outcome of RD&D investments is characterized using novel expert elicitation data, allowing for a transparent and consistent integration into the framework. The problem is formulated as a multi-stage decision under uncertainty to represent opportunities for policymakers to learn and adapt to new information between decision stages. Results show that under a carbon constraint, the optimal investment strategy includes lower solar PV RD&D spending upfront but more RD&D spending later—and sometimes higher spending overall—when compared to a strategy under perfect foresight about RD&D outcomes, or based on single-shot decision-making under uncertainty without learning. We also show that when uncertainty is considered without learning, new solar PV capacity investments are depressed. Overall, the results caution that the most common approaches used by the research community today to inform policy on optimal energy technology investments, including scenario analysis and Monte Carlo simulation that assume perfect foresight with no learning, may be systematically over- or under-estimating the optimal investment strategy.

Keywords

Electricity sector; technology investment planning; R&D; decision-making under uncertainty; learning; solar

Acknowledgements

The authors wish to thank William Hogan, Gabe Chan, and Joern Huenteler for comments on an earlier draft, and Mort Webster and Jennifer Morris for helpful discussions about stochastic dynamic programming for energy and climate policy modeling. They also gratefully acknowledge financial support from the Harvard Belfer Center Science, Technology, and Public Policy program and Energy Technology Innovation Policy research group, Shell, and BP.

Table of Contents

1. Introduction

Policymakers and other stakeholders currently face the tremendous challenge of addressing the role future electricity systems will play in realizing long-term climate change mitigation goals. Of particular urgency is a need to balance near-term decisions about ongoing investments in new, relatively expensive, low-carbon electric power generating capacity, with investments in research, development, and demonstration (RD&D) aimed at innovating and commercializing the next generation of competitive, low-carbon generation technologies (DOE 2011; Nemet and Baker 2009). Achieving cost-effectiveness in promoting future low-carbon electricity systems, however, requires adaptively managing inherent uncertainties about the cost and availability of emerging technologies (NRC 2007).

Solar photovoltaic (PV) technology is at the forefront of the U.S. national energy policymaking agenda, and highlights this challenge particularly well (DOE 2012). On one hand, new generating capacity investments—which may include commercially available solar PV—must continue to maintain a reliable electricity supply. On the other hand, ongoing public solar PV RD&D investments can help lower the overall cost of managing carbon emissions and mitigating climate change if the RD&D return rate is favorable. Given (1) the long-term strategic nature of the energy investment and climate change problem; (2) the inherent uncertainty in the outcomes to RD&D; (3) the opportunity to learn from interim investments and adapt future decisions; and (4) the complementary *and* competing roles that different technologies play within the power system, what tools can best support policymakers and other stakeholders in determining an optimal balance of technology deployment and solar PV RD&D investments over time?

Current decision support tools for such "deployment versus development" questions in the electricity sector remain restricted in scale or scope. First, most models that contain a resolution of the physical electric power system sufficient to model its unique characteristics and technology interdependencies (e.g., temporal load variability, plant operations and constraints, energy demand balance, supply reliability, resource availability), *and* that contain an endogenous representation of the effect of research on technology costs, still represent decisions as being made up front as "single shots" assuming perfect information (e.g., Barreto and Kypreos 2004; Kouvaritakis, Soria, and Isoard 2000). This shortfall remains burdensome given evidence that the optimal decision in complex systems often diverges under uncertainty and learning (e.g., Webster 2002; Webster, Santen, and Parpas 2012).

When uncertainty has been incorporated, computational limitations have prevented it from being appropriately considered in the context of multi-stage decision-making and interim learning. Onerously, such problems are particularly prone to the "curse of dimensionality," where their size grows exponentially large with the number of decisions, uncertainties, and decision periods normally considered in studies of real-world problems (Powell 2011). Moreover, this curse of dimensionality is exacerbated when the uncertainty in the problem is either decision-dependent or continuous, or both. Thus, almost 35 years after Manne and Richels (1978) presented a sequential

stochastic decision analytic framework for energy technology R&D decisions[1], deterministic models are still widely used in combination with scenario analysis or formal Monte Carlo analysis as the method to integrate uncertainty, missing the opportunity to learn and adapt (e.g., Bergerson and Lave 2007; Richels and Blanford 2008). More complex methods of accounting for the uncertain nature of research outcomes have also been employed in the context of electricity capacity and R&D investment planning through mapping alternative R&D-dependent technology pathways to outcomes of a deterministic model on an expected value basis (e.g., Blanford 2009), or introducing expected value-based penalties into the objective cost function of models (e.g., Messner et al. 1995; Grubb 2002; Grubler and Gritsevksii 1997). However, neither method considers the realistic ability to adapt decisions over time because the impacts of early R&D investments are still assumed at the outset. Finally, stochastic programming methods with recourse have been used with a few models to introduce adaptive decisions and learning in a formal stochastic structure (e.g., Bosetti and Tavoni 2009; Kypreos et al. 2000). Unfortunately, the use of exogenous scenario trees to represent uncertainty in stochastic programming still makes the approach relatively susceptible to dimensionality burdens. This feature of stochastic programming is particularly salient in the type of decision-dependent uncertainty inherent in the energy innovation process, where decisions about RD&D (e.g., level of spending, timing of investments) can change future return likelihoods by pushing the research frontier outward (Popp, Newell, and Jaffe 2009).

Second, existing datasets linking research investments to improvements in technology costs have been relatively disconnected from the models that can support decisions about optimal electricity capacity and RD&D investments under uncertainty. This has challenged the integration of consistent and transparent estimates of the uncertainty in the returns to RD&D. For example, empirical studies using innovation "inputs" (e.g., RD&D expenditures) and installed capacity to explain historical cost-reductions are often limited by the fact that improvement can be affected by other factors, such as economies of scope and material prices. Additionally, these studies do not typically account for the heterogeneous nature of RD&D investments, or the poor availability of private sector RD&D data (Popp, Newell, and Jaffe 2009; Qiu and Anadon 2012). Other studies using innovation "outputs," such as patents and patent citations to study the effect of RD&D on technical change (e.g., Griliches 1990; Popp 2002; Popp et al. 2013) address some of these shortfalls, but still suffer from the fact that using historical information to model future innovation may be inconsistent with reality (Chan et al. 2011). Additionally, the format of results from patent and citation studies are often hard to translate into units of technical change that can be transparently introduced into detailed "bottom-up" electricity models. Thus, many models still rely on stylized assumptions about research returns (e.g., Bosetti and Tavoni 2009; Blanford 2009); others that integrate more consistent data do so from a "top-down" economic perspective that limits the ability to study R&D effects in the context of electric power system-level generation capacity expansion (e.g., Popp 2004; Popp 2006).

Ultimately, most existing decision-support tools are either (1) still mismatched with the manner in which long-term strategic investment decisions for the electricity sector are actually made—at multiple times throughout the planning horizon, after interim collection and assessment of new information, and in the context of multiple interacting technologies within a power system, or (2)

[1] Manne and Richels (1978) solve for optimal U.S. breeder reactor R&D strategies using a probabilistic decision-tree with a small number of possible decisions (four), and a highly discretized uncertainty space (e.g., "yes/no," high/medium/low" scenarios). The problem, as simplified, does not therefore suffer from the curse of dimensionality.

based on stylized or inconsistent estimates about uncertainty in future technology costs. These issues can result in improper decision support to policymakers, suggesting investment strategies that are either too high or too low, or incorrectly timed and targeted, when compared to the optimal, cost-effective strategy under realistic uncertainty and the opportunity to learn.

In this paper, we present a new bottom-up, stochastic modeling framework with endogenous technical change for supporting electricity generation capacity and solar PV RD&D investment decisions under decision-dependent uncertainty and learning. The new model is formulated as a formal sequential decision under uncertainty problem, and solved using multi-stage stochastic dynamic programming. Specifically, we adapt emerging approximate dynamic programming techniques to manage the dimensionality burden that arises from integrating continuous uncertainties through multiple decision periods, and the complexity of representing the decision-dependent uncertainties reflective of the innovation process. In a stepwise manner, we show the value of using the new stochastic approach over dominant "single-shot" approaches, which do not integrate uncertainty, and scenario approaches, which account for uncertainty but do not account for the opportunities policymakers have in adapting their investment decisions over time. We also characterize the uncertainty in the solar PV RD&D process using expert elicitation data, which allows for a transparent and consistent integration into the bottom-up modeling framework. Results from the new framework can inform policymakers about how to balance near-term spending on deployment versus development programs for emerging low-carbon electricity technologies in the face of RD&D uncertainties.

The remainder of the paper is organized as follows. Section 2 presents the new modeling approach, outlining the multi-stage decision problem and characterization of the solar PV RD&D uncertainty. Section 3 presents optimal investment strategy results from the stepwise comparison of a single-shot deterministic approach, a formal Monte Carlo analysis to integrate uncertainty, and the formal stochastic dynamic programming approach to integrate both uncertainty and learning, and does so within the context of two different carbon policy regimes. Section 3 also presents results of a sensitivity test to check robustness of the results. Section 4 concludes with summary discussion and suggestions for future research.

2. Methods

2.1 Model Overview

We formulate the research question as a multi-period sequential decision under uncertainty problem, with an opportunity to learn and adapt decisions between decision stages. The model consists of four decision stages, with decadal time steps intended to cover approximately 2010-2050. New power plant capacity investment decisions for four possible technology groups (conventional coal, conventional natural gas, wind, and solar PV) and RD&D investment decisions for carbon-free solar PV are made in each decision period to operate the power system and reliably meet electricity demand, and to propagate additional technical change in solar PV, *prior* to learning about the amount of technical change the RD&D will produce. RD&D investment is represented as a continuous range of values between "business-as-usual" U.S. public spending on solar PV and

ten times the amount of RD&D recommended by technology experts in the field (Anadon et al. 2011)[2]. After each decision stage, the decision-maker learns about the realized amount of technical change, and has an opportunity to adapt her next decision based on the "state of the world" she finds herself in. The decision-maker in this context is a hypothetical central decision-maker who seeks to minimize total system costs for both capacity deployment and emerging technology development[3]. As described in more detail below, uncertainty about solar PV innovation is modeled as a decision-dependent process where higher levels of RD&D correspond to a greater likelihood of higher future returns (i.e., larger investment cost reductions). Operation of the underlying power system and reliability of supply are represented using a least-cost generation capacity expansion planning approach, with optimal dispatch. In each case, a cumulative carbon emissions cap is also in place to simulate a climate change policy (or no cap in the case of no policy).

We use an underlying electricity system that consists of a large base of conventional pulverized coal and natural gas plants, some wind power, and a small base of emerging utility-scale solar PV. Combined, the total gigawatts installed in the base year approximates the size of the U.S. system in 2010 (EIA 2010). The base year for electricity load is also an approximation of annual U.S. electricity demand, specified in seventeen segments for time of day and season. Full model details, parameter values, and underlying electricity system data are given in an appendix, available upon request.

2.2 Mathematical Formulation

The objective of the stochastic problem is:

$$\min_{NC_{g,1},RDD_{g*,1}} \left\{ C_1 + \min_{NC_{g,2}RDD_{g*,2}} E_{\theta_{g*,t}}[C_2 + \cdots] \right\},$$

where NC is new capacity investment, RDD is the RD&D investment level in solar PV, θ is the uncertain technical change for solar PV, C is the discounted total system cost, and subscripts g, $g*$, and t represent technology groups, emerging solar PV, and time, respectively. Stochastic dynamic programming is used to structure the problem, an approach that uses the Bellman equation (Bellman 2003) to decompose the multi-period problem into a simpler set of conditions that must hold for all decision stages:

$$V_t = \min_{NC_{g,t}\, RDD_{g*,t}} \left\{ C_t + E\left[V_{t+1}\left(NC_{g,t}, RDD_{g*,t}, \theta_{g*,t} \right) \right] \right\},$$

where V_t is the "Bellman Value."

[2] 2010 BAU solar PV spending = $143M, the recommended amount of solar PV spending by the median technology expert = $200M, and ten times the median recommended amount = $2000M (Anadon et al. 2011)

[3] For simplicity, we assume a perfect market and can thus consider RD&D investments from the perspective of the public RD&D decision-maker, and capacity investments from the perspective of the utility or other power producer.

Costs in each stage are computed as:

$$C_t = \sum_{g,g*} \left(OC_{g,t} + FC_{g,t} + VC_{g,t} + SV_t + RDD_{g*,t} \right) \cdot (1+r)^{-t},$$

where *OC* is the total overnight capital cost for new power plant installations, *FC* are fixed operation and maintenance costs, *VC* is the variable (and fuel) cost from optimally dispatching the technologies to meet electricity demand in each of seventeen time-of-day and seasonal load duration curve slices, *SV* is a potential penalty for breaking the carbon cap (equivalent to executing a carbon emissions "safety valve" option), *RDD* is the solar PV RD&D investment; *r* is the discount rate[4].

2.3 *Uncertainty in solar PV RD&D-induced technical change*

We explicitly represent the innovation process and the uncertainty in returns to solar PV technology research in the modeling framework. Specifically, we model the type of innovation propagated by public RD&D spending.[5] This technical change enters the model as explicit RD&D-induced overnight capital cost reductions in utility-scale solar PV technology, via a shifted log-logistically distributed random variable conditional on the level of RD&D:

$$oc_{g*,t+1} = oc_{g*,t} - \theta_{g*,t},$$

where *oc* is the overnight capital cost per GW solar PV, θ is its realized cost reduction, and μ, σ, and ξ are the location, scale, and shift parameters of the shifted log-logistic density function.

We use the solar PV capital cost dataset resulting from a recent large-scale expert elicitation study that gathered estimates about future costs and performances of several emerging energy technologies, including their uncertainties, for various levels of RD&D (Anadon et al. 2011), and develop shifted log-logistic probability density functions for the RD&D-cost reduction relationships by extending the method outlined by Chan and Anadon (2014). We employ conditional density functions to represent the feature of decision-dependence inherent in the innovation process. Full details on the construction of the density functions and data used, as well as examples of the resulting functions, are provided in an appendix, available upon request.

[4] Note that we use a centralized electricity market structure for the generation capacity expansion and optimal dispatch problem. We use this over a decentralized representation due to the long-term strategic nature of the investments under study in this research, and the large structural uncertainties present in short-term market behaviors of individual firms over such long time frames (Pérez-Arriaga and Meseguer 1997)

[5] Historical public U.S. RD&D spending data for each of the emerging technology groups was provided to the technology experts as background information during the expert elicitation study, and experts were asked to base their cost and uncertainty estimates on public RD&D spending and its possible effects (Anadon et al. 2011).

2.4 Solution Methods

The stochastic dynamic programming (SDP) sequential decision problem formulated above has a finite horizon, and is traditionally solved as a Markov Decision Problem (Bertsekas 2007) using a backward induction algorithm. However, the algorithm exhaustively iterates over the state, decision, and uncertainty spaces for each decision period to calculate the exact Bellman Value function and corresponding "policy function" (decision strategy) in each stage. Due to the continuous nature of the capacity and RD&D investment decisions in this problem, as well as the continuous and decision-dependent nature of the uncertainty, we employ approximate dynamic programming (ADP) techniques to manage the curse of dimensionality encountered, and solve the problem more efficiently. ADP is a family of methods (e.g., Bertsekas and Tsitsiklis 1996; Powell 2011) that approximates the value function in each period by adaptively sampling the state space to focus on lower expected value states until the Bellman Value function converges. We use an ADP algorithm developed by Webster, Santen, and Parpas (2011), and employed by Santen (2012) for electricity sector R&D investment planning under uncertainty, which uses a two-phased approach of constrained Latin Hypercube Sampling followed by random sampling to search over the state space; and a Moving Least Squares method to approximate future Bellman Values via interpolation.

For the comparison of methods exercise, we also constructed a deterministic non-linear programming (NLP) version of the model with underlying features identical to the SDP. The NLP assumes perfect foresight about returns to RD&D. We use the deterministic model directly to study the decision problem under no uncertainty and no learning, and in a Monte Carlo analysis to study the decision problem under uncertainty but no learning. We use 500 modified Latin Hypercube Monte Carlo samples of solar PV capital cost reduction uncertainty and the deterministic version of the model to perform the Monte Carlo analysis.

3. Results

In this section, we present results from solving the electricity capacity and RD&D investment decision problem with and without uncertainty and learning, and under two different carbon policy regimes. We present the results from each subsequent case (first the stochastic dynamic program, then the Monte Carlo analysis) against the original deterministic results in order to explicitly compare the effect uncertainty and learning has on the investment strategy, with the effect state of the art methods of incorporating uncertainty (but *no* learning) through Monte Carlo analyses have. The two carbon policy regimes explored include a "no policy" case with unconstrained carbon emissions, and a policy case with a moderately stringent cumulative carbon emissions cap equivalent to reducing annual emissions from the power sector fifty percent from "business as usual" in 2010. Finally, we present and discuss robustness of the results from a sensitivity analysis on RD&D-based technical change using additional expert elicitation data from Anadon et al. (2011).

3.1 Optimal investment strategy under uncertainty and learning

When no carbon policy is present and electricity sector emissions are allowed to grow unconstrained over time, the investment strategy under uncertainty and learning differs only negligibly from the deterministic strategy (Figures 1 and 2). In both cases, the capital investment strategy involves investing in new coal and gas plants, and the near-term solar PV RD&D investment strategy differs only by about one percent of the maximum possible RD&D investment level. This is expected given the lack of incentive to reduce carbon emissions from business-as-usual[6]. The minor difference in the RD&D strategy is anticipated, due to use of the heuristic-based ADP method that has a small degree of irreducible error (Powell 2011). Note that consistent with the output of solving stochastic decision problems with learning, we present the first period decision as a single implementable decision (Figure 1), and future period decisions as a "decision rule" conditional on the state of the world (Figure 2).[7]

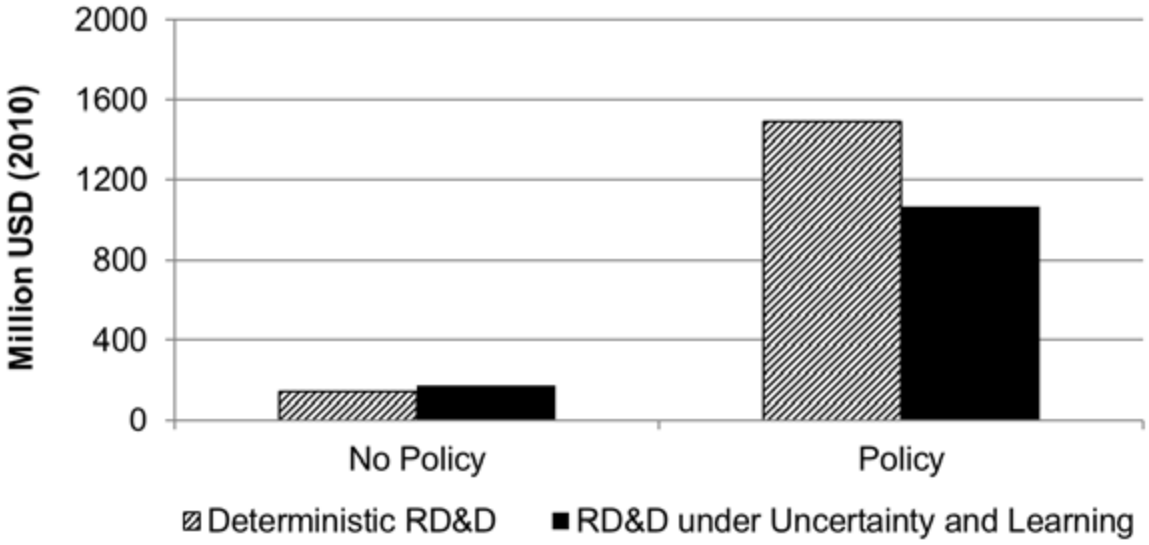

Figure 1 Optimal first period annual solar PV RD&D under uncertainty and learning (black bars) and under perfect foresight (hashed bars) with no carbon policy (left) and with a carbon policy (right)

In the presence of a moderate carbon constraint the optimal solar PV RD&D investment strategy diverges considerably from the deterministic investment strategy. Overall, the carbon constraint favors deploying renewable technologies such as wind and solar (over coal and gas) to reduce carbon emissions, and developing solar PV through RD&D to reduce costs. However, under uncertainty and learning, the near-term (first-period) optimal solar PV RD&D investment is substantially lower than the deterministic strategy (although still above BAU spending levels) (Figure 1). This change represents a difference of over 20 percent of the maximum possible RD&D investment level. In contrast, Figure 2 compares the range of optimal RD&D investment

[6] For simplicity, and to isolate the effect of carbon policy on technology choice, we ignore other potential drivers of solar PV deployment such as high fossil fuel costs or competing technology capacity constraints.

[7] The minimum level of RD&D, representing business-as-usual, is always optimal for the third and fourth periods due to the temporal nature of the decision problem. No new capacity investments are made in the fourth period due to an "end of the world" effect after the last period, and it is never optimal to invest more than the minimum investment in these later periods because RD&D investments affect capital costs in the *next* decade.

decisions in the stochastic decision rule and the deterministic optimum, showing that the optimal RD&D investment under uncertainty and learning is *higher* than under perfect foresight for all possible technical change futures in the second period. Under some realizations of the future, optimal RD&D spending in the second period is so high that even cumulative RD&D spending is higher under uncertainty and learning than under an assumption of perfect foresight.

(a) No carbon policy case

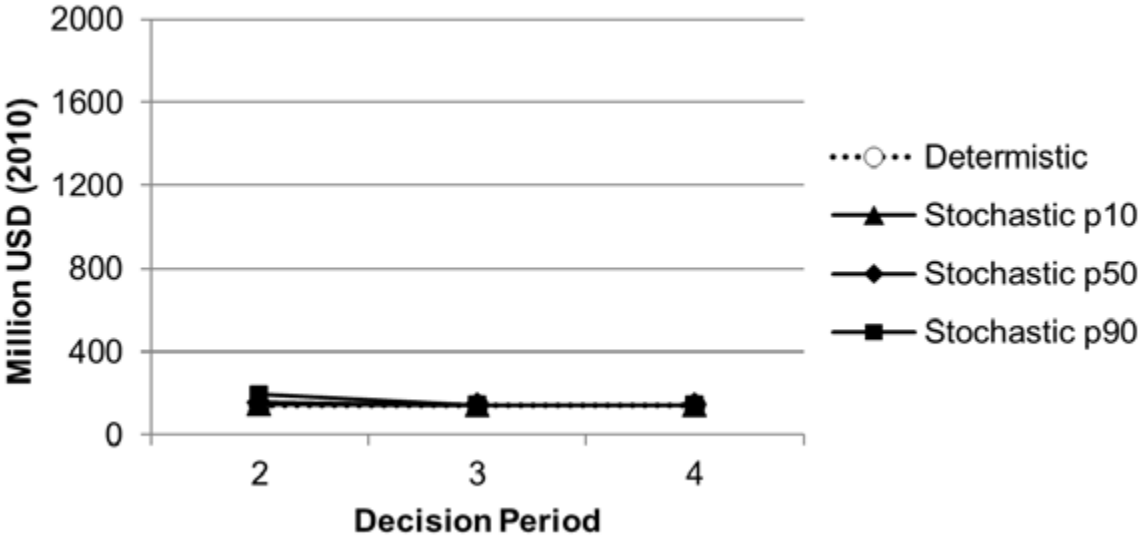

(b) Carbon policy case

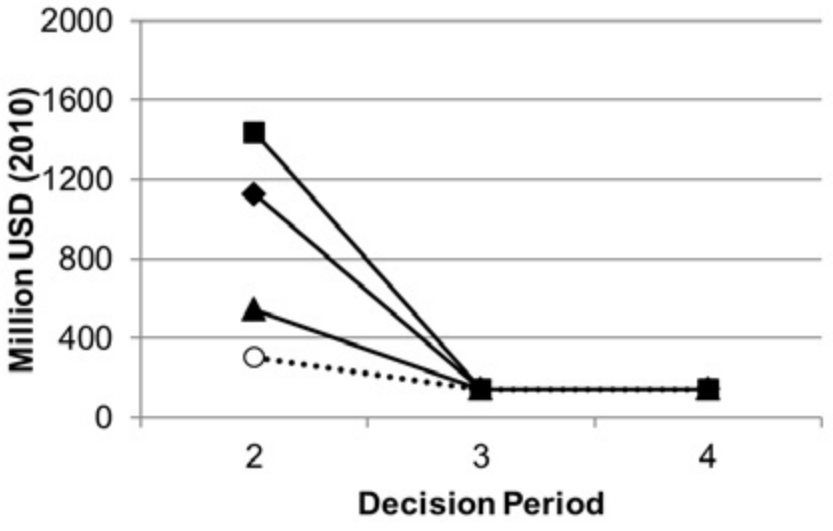

Figure 2 Optimal future annual solar RD&D investments under uncertainty and learning (solid lines) and under perfect foresight (dashed line) (Note: The stochastic investment strategy is presented as a decision rule with percentiles because the optimal future decision is conditional on the future state)

The optimal solar PV RD&D investment is lower in the near term than the deterministic strategy so that the benefits of learning between Period 1 and Period 3 (when heavy solar PV deployment occurs to meet the cumulative carbon cap) may be fully realized. This pattern of waiting until later periods to meet a cumulative emissions cap with deployment of low-carbon technologies such as solar that are farther from the market, and investment in their RD&D upfront, is consistent with the findings in the electricity R&D investment planning model of Santen (2012).

It is noteworthy that some positive RD&D over the minimum BAU level is optimal in the first period in the stochastic case. One might ask, if waiting to learn is the best strategy, why invest anything at all upfront? The decision-dependent nature of the RD&D uncertainty in our work helps explains this result. Investing early in RD&D provides a snowball effect on reducing the cost of solar PV by the time it needs to be built. The higher RD&D investment occurs in order to take advantage of potential early "tail" opportunities for return on the investment, and propagates this snowball effect. Meanwhile, if early realized RD&D returns do not turn out to be favorable (only modest cost reductions are realized), there is still time before solar PV deployment to adapt the investment strategy.

Lastly, once again there is negligible difference between the stochastic and deterministic results for new capacity deployments. Given the specific capital cost reduction potentials considered, there remains a dominant deployment plan that meets all constraints in a least-cost manner. In the presence of the carbon constraint, RD&D investment alternatives afford an opportunity to reduce total system costs further.

3.2 Optimal investment strategy under uncertainty without learning

When RD&D uncertainty continues to be considered, but the realistic opportunity for the decision-maker to learn about solar RD&D returns between decision periods and adapt her decisions is still ignored, the optimal strategy diverges from the strategy under perfect foresight in a different way. Under no carbon policy, considering uncertainty changes neither the optimal RD&D investment path nor the capacity investment path from when perfect foresight is assumed. As in the previous section, solar PV is not part of the optimal deployment plan when emissions are unconstrained; considering a known median level of return on RD&D investments, the decision-maker knows to choose the minimum level of RD&D possible.

Under a carbon policy, considering uncertainty reduces the level of solar PV RD&D investment in both the first and second period from the deterministic strategy (although once again it is still higher than current BAU spending levels) (Figure 3). Note that in the case of decision problems without learning, for which a deterministically structured model is used in a Monte Carlo analysis, the result is the familiar optimal implementable investment path (in contrast to a decision rule). In this case, we use mean outputs for capacity and RD&D investments to compute the optimal path under uncertainty.[8]

[8] Resulting optimal new capacity investment paths were re-tested in the optimal dispatch component of the model to ensure that all electricity system constraints remain satisfied and that the investment strategy was viable.

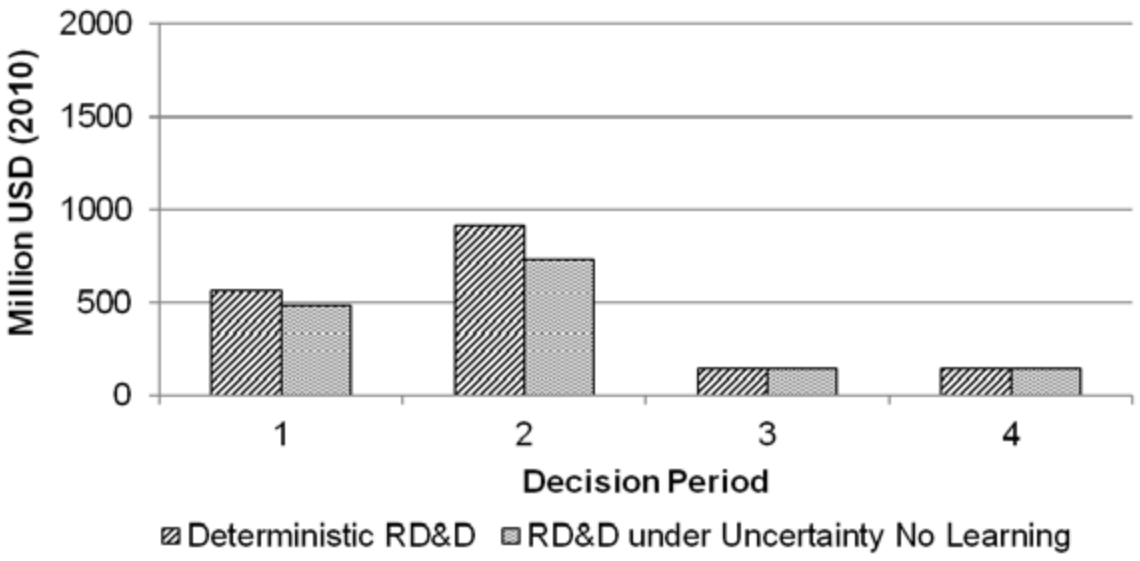

Figure 3 Optimal annual solar RD&D with uncertainty (zigzagged bars) and without uncertainty (hashed bars) under a carbon policy

The decrease in first and second period solar PV RD&D investments can be explained by the associated capacity investment strategy. While there is negligible difference across first period capacity investments under uncertainty and no uncertainty (Figure 4a), there is a large difference between investments by the third period (Figure 4b). Under uncertainty, the capacity investment strategy involves investing twice as much wind as solar PV, whereas under a perfect foresight assumption this balance flips and the strategy involves investing almost twice as much solar PV as wind capacity. In this case of uncertainty without learning, it is the lower level of solar PV deployment under uncertainty that produces, on average, an RD&D investment strategy that remains lower than under perfect foresight throughout the planning horizon. Overall, ignoring RD&D return uncertainty underestimates the role of the emerging technology in the physical system (i.e., underbuilds solar PV), underestimates the amount RD&D investment in the first period should be lower, and incorrectly assumes that continue lower investment is optimal when RD&D should actually be higher in the second period.

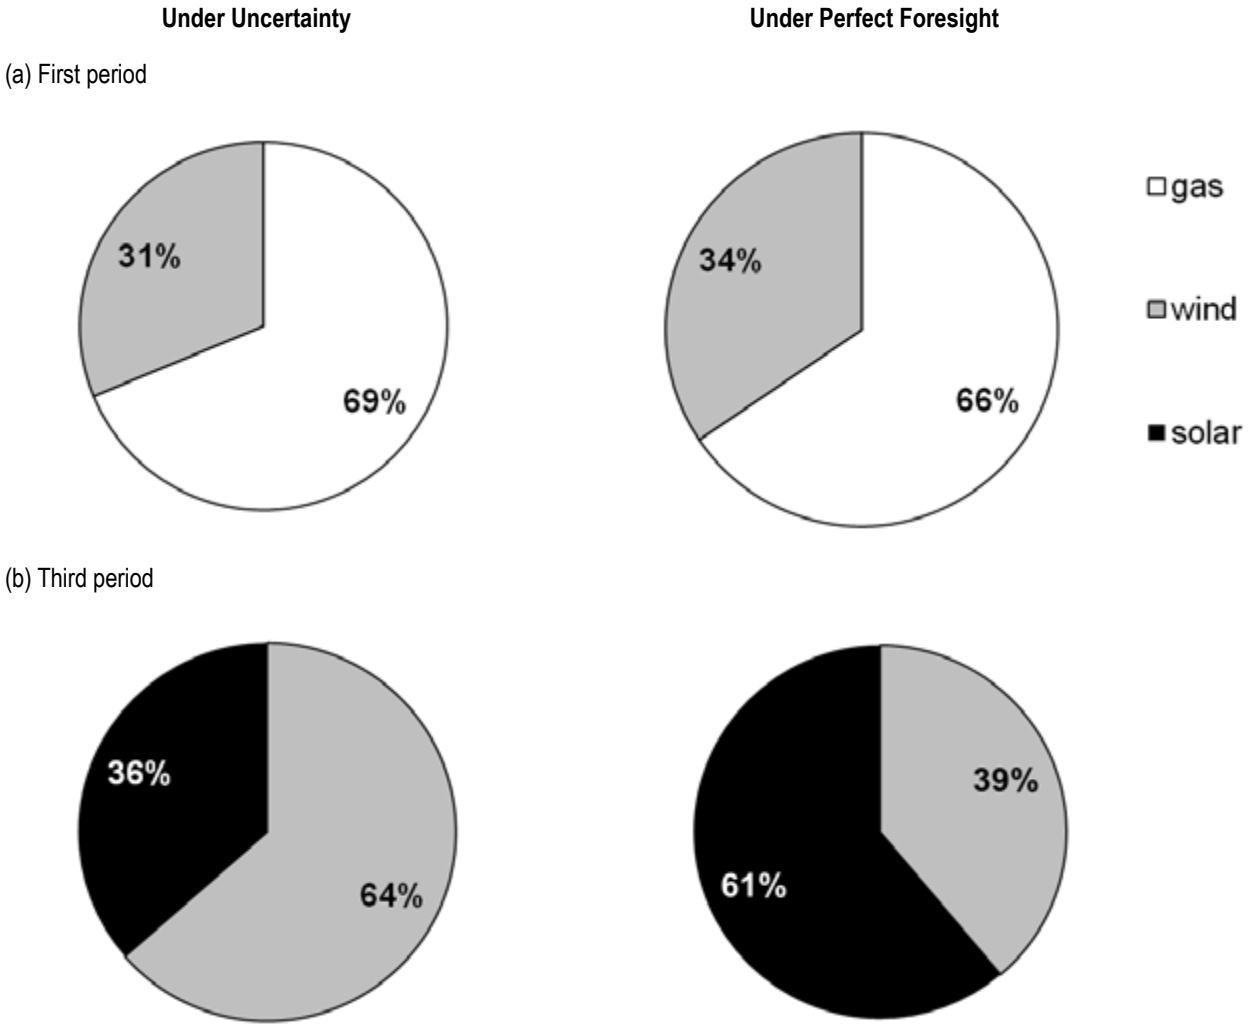

Figure 4 Optimal capacity investment ratios under uncertainty (left) and under perfect foresight (right); Legend: gas = white; wind = gray; solar = black; there is no new coal capacity in these scenarios

3.3 Sensitivity to technical change uncertainty

We check the robustness of the results to the level and characterization of the uncertainty in RD&D returns. The expert elicitation dataset of Anadon et al. (2011) employed contains future overnight capital cost estimates for solar PV from a variety of solar technology experts, and the original stochastic model uses estimates based on the "median" expert. In this subsection, we present results from solving the stochastic model with uncertainty characterized instead using the most optimistic and most pessimistic experts, and compare the optimal investment strategy to the original model's optimal investment strategy, as well as to their respective deterministic investment strategies. "Optimistic" corresponds to the expert that estimated the 2030 costs of utility solar PV to be the lowest, while "pessimistic" corresponds to the expert that estimated 2030 costs

to be the highest. Online Resource 3 shows examples of the resulting alternate probability density functions and provides details about the data used for their construction. We apply a moderately stringent carbon policy for this discussion, as used above.

Overall results from the sensitivity analysis show that lower near-term RD&D investment, higher second period RD&D investment, and a relatively fixed capacity deployment plan (compared to using an assumption of perfect foresight), are robust to a wide range of future cost possibilities for utility solar PV. We also show a trend of increasing divergence of the optimal first-period RD&D investment strategy under uncertainty and learning from the deterministic strategy (i.e., greater levels of pessimism about future cost reduction possibilities induce additional lowering of first-period RD&D from the deterministic strategy on a percentage basis) (Figure 5). Moreover, very optimistic future costs for solar PV allow the stochastic solution to begin approaching the deterministic solution (in Figure 5, the change is only modest, at less than 1% of the maximum optimistic RD&D level). The optimal RD&D investment strategy continues to be higher than BAU spending levels throughout the sensitivity analysis, as well.

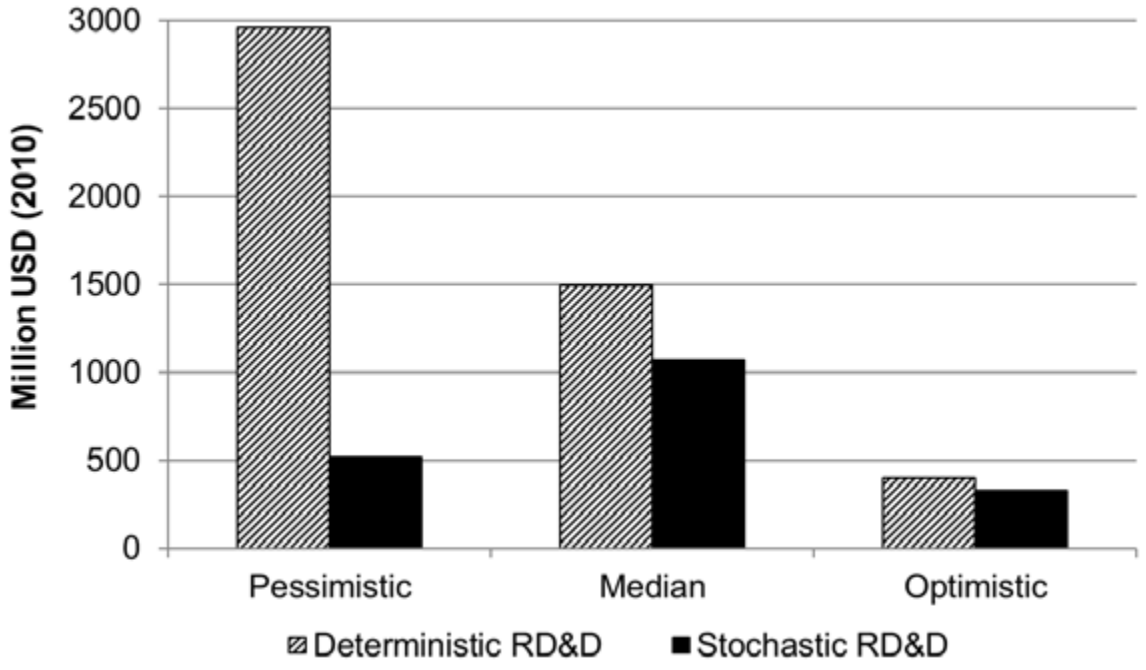

Figure 5 Optimal first-period RD&D investments under three RD&D return uncertainty estimates and learning. (Note: Pessimistic RD&D Max = $3000M per year; Median RD&D Max = $2000M per year; Optimistic RD&D Max = $10,000M per year)

4. Concluding Discussion

In this paper, we presented a new modeling framework for simultaneously choosing sequential capacity and solar PV RD&D investments for the electric power generation sector under decision-dependent technical change uncertainty and carbon constraints. We numerically solved for the optimal investment strategy under uncertainty and learning using stochastic dynamic programming, and compared the optimal strategy to numerical solutions of the decision problem under uncertainty and no learning, as well as under perfect foresight (deterministic with no learning). Through this explicit comparison, we showed the importance of using a formal stochastic approach over the current dominant approach of using scenario analysis or Monte Carlo analysis to study optimal investment strategies under uncertainty in this context.

Results show that under a moderately stringent carbon policy, the optimal RD&D investment strategy for an emerging low-carbon technology such as solar PV involves spending less upfront when uncertainty and learning are explicitly considered than when perfect foresight is assumed, and involves adapting the investment decision to new information in the second period by increasing RD&D investment accordingly. This is true even though it is not optimal to change the capacity deployment plan in the presence of uncertainty and learning. We also show that, despite the first period investment being lower and second period investment being higher than the deterministic strategy, 1) RD&D spending under a carbon policy is always higher than current BAU spending levels, and 2) under some possible technical change futures, the second period RD&D investment can be high enough to result in higher RD&D spending overall.

Moreover, we revealed that a Monte Carlo analysis of the same decision problem, which considers uncertainty but no learning, underestimates the amount solar PV RD&D investment should be lowered from the deterministic strategy, and incorrectly indicates that keeping RD&D investment lower in the second period is optimal. With respect to capacity deployment, the Monte Carlo analysis also incorrectly suggests that in the presence of possible R&D failures, capacity investment switching is optimal. In fact, with the opportunity to learn and adapt the RD&D investment decision over time, the stochastic model shows that it is not always necessary to change the original deployment plan. Figure 6 summarizes the first-period RD&D investments across the three cases presented in this paper by showing the optimal investment strategies normalized as a percentage change from the maximum RD&D investment possible. Under no policy, the differences between optimal solutions are negligible and ultimately within error bounds of the ADP method. However, in the presence of a carbon policy when there is additional solar PV deployment, we caution that over-reliance on Monte Carlo or other scenario methods using deterministically structured models may be providing improper decision support in this context. These methods do not capture 1) the full extent of the optimal RD&D reduction from the deterministic strategy in near-term; 2) that RD&D investment in the second period is actually be *higher* than the deterministic strategy, not lower; and 3) the fact that capacity switching is not always optimal.

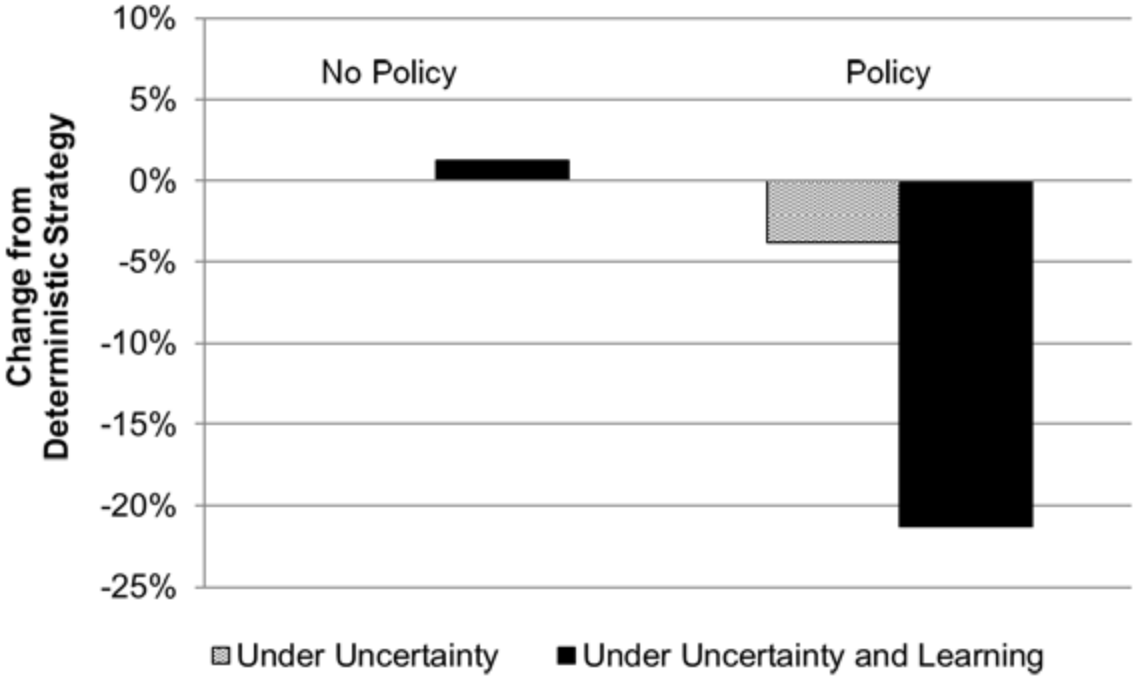

Figure 6 Change in first-period solar PV RD&D investment strategies from the deterministic strategy under no carbon policy (left) and a carbon policy (right), with respect to maximum possible RD&D investments. Legend: RD&D under uncertainty (no learning) = zig-zagged bars, stochastic RD&D = black bars (Note that under "no policy," the optimal strategy under uncertainty (no learning) is the same as the deterministic strategy and therefore has a value of zero on this percentage change graph)

The ability to learn and incorporate new information between decision periods is a key feature producing our results. In the carbon policy case under uncertainty but no learning, the optimal capacity investment strategy for solar PV was lower than the optimal strategy under perfect foresight. Solar PV RD&D investment decreased as a natural result (there was less overall benefit from the investment). However, under uncertainty and learning, the optimal capacity investment strategy for solar PV remained the same as the strategy under perfect foresight. Why, then, did the near-term RD&D investment decrease even further? With learning and adaptation explicitly accounted for in the modeling framework, the optimal near-term strategy incorporated future benefits from risky RD&D investments and high-reward "tail event" cost reductions by expanding the decision rule in future periods. Likewise, the impact of downsides from high-risk RD&D investments that do not yield favorable returns can be reduced because in the case of observing an unfavorable RD&D outcome, the decision-maker can select a new investment path to help her stay as close to the global optimum as possible. When this flexibility is possible and the decision-maker need not select a single path at the beginning, she can "hedge" against future uncertainties by investing more modestly at the start. Higher RD&D investment in the future can also be discounted, lowering the net present value of the total investment costs.

The specific nature of the uncertainty incorporated also played an important role in constructing the optimal investment strategy under uncertainty and learning. Using alternate data on RD&D-based solar PV cost-reduction potentials from the most optimistic and pessimistic solar technology experts in Anadon et al. (2011), we showed that under a carbon policy the magnitude of the first-period hedge varies substantially, although it remains optimal to keep RD&D investment lower than the deterministic strategy. When it is possible to learn and adapt to interim information about the RD&D return rate, incremental RD&D steps are optimal. Furthermore, the sensitivity analysis suggests that increasing pessimism about the future costs of solar PV only amplifies the notion of ramping RD&D investment up after learning can occur.

Future research should focus on scaling the stochastic modeling framework outlined in this paper to an industrial-scale energy systems model with additional electricity technologies, RD&D choices and associated RD&D uncertainties, and uncertainties such as fuel price and renewable resource availability. With respect to capacity and RD&D decisions, it would be helpful to compare the effect of multiple interacting technologies competing for RD&D investment with the solar-only RD&D program studied in this paper, to further comment on the generalizability of results. Additionally, it would be valuable to develop frameworks for integrating innovation on demand-side technologies and electricity storage opportunities. Including technical change in emerging technologies such as plug-in hybrid and electric vehicles, automated demand response programs, smart grids, and electricity storage would complement the existing supply side technical change outlined in the framework.

References

Anadon LD, Bunn M, Chan G, Chan M, Jones C, Kempener R, Lee A, Logar N, and Narayanamurti V (2011) Transforming U.S. Energy Innovation. Belfer Center for Science and International Affairs, Harvard Kennedy School, Cambridge

Barreto L, Kypreos S (2004) Endogenizing R&D and market experience in the "bottom-up" energy systems ERIS model. Technovation 24: 615-629

Bellman R (2003) Dynamic programming. Dover Publications, New York

Bergerson JA and Lave LB (2007) Baseload coal investment decisions under uncertain carbon legislation. Environ Sci Technol 41: 3431-3436

Bertsekas D (2007) Dynamic Programming and Optimal Control Vol I and II 3rd edn. Athena Scientific, Nashua

Bertsekas D, Tsitsiklis J (1996) Neuro dynamic programming. Athena Scientific, Belmont

Blanford GJ (2009) R&D investment strategy for climate change. Energy Econ 31: S27-S36

Bosetti V, Tavoni M (2009) Uncertain R&D, backstop technology, and GHGs stabilization. Energy Econ 31: S18-S26

Chan G, Anadon LD, Chan M, Lee A (2011) Expert elicitation of cost, performance, and RD&D budgets for coal power with CCS. Energy Procedia 4: 2685-92

Energy Information Administration (EIA) (2010) Electric Power Annual 2010. http://www.eia.gov/electricity/annual/. Accessed 22 October 2013

Griliches Z (1990) Patent statistics as economic indicators: a survey. J Econ Lit 28: 1661-1707

Grubb M, Kohler J, Anderson D (2002) Induced technical change in energy and environmental modeling; analytic approaches and policy implications. Annu Rev Energy Environ 27: 271-308

Grubler A, Gritsevskii A (1997) A model of endogenous technological change through uncertain returns on learning (R&D and investments). Conf Induc Technol Change Environ

Kouvaritakis N, Soria A, Isoard S (2000) Modelling energy technology dynamics: methodology for adaptive expectations modeling with learning by doing and learning by searching. Int J Global Energy Issues 14: 104-115

Kypreos S, Barreto L (2000) ERIS: a model prototype with endogenous technological change. Int J Global Energy Issues 14: 374-396

Manne AS, Richels RG (1978) A decision analysis of the U.S. breeder reactor program. Energy 3: 747-767

Messner S, Golodnikov A, Gritsevskii A (1995) A stochastic version of the dynamic linear programming model MESSAGE III. Energy 21: 775-784

National Research Council (NRC) (2007) Prospective Evaluation of Applied Energy Research and Development at DOE (Phase Two). The National Academies Press, Washington. http://www.nap.edu/catalog.php?record_id=11806. Accessed 11 November 2013

Nemet G, Baker E (2009) Demand subsidies versus R&D: comparing the uncertain impacts of policy on a pre-commercial low-carbon energy technology. Energy J 30: 49-80

Pérez-Arriaga, IJ, Meseguer C (1997) Wholesale marginal prices in competitive generation markets. Power Systems, IEEE Transactions on 12: 710-717

Popp D (2002) Induced innovation and energy prices. Am Econ Rev 92: 160-180

Popp D (2006) They don't invent them like they used to: an examination of energy patents over time. Econ Innov New Techn 15: 753–776

Popp D, Newell R, Jaffe A (2009) Energy, the environment, and technological change. NBER Working Paper 14832. http://www.nber.org/papers/w14832. Accessed 17 January 2014

Popp D, Santen N, Fisher-Vanden K, Webster M (2012) Technology variation vs. R&D uncertainty: what matters most for energy patent success? Resour and Energy Econ 35: 505-533

Powell WB (2011) Approximate dynamic programming: solving the curses of dimensionality. John Wiley & Sons, Hoboken

Qiu Y, Anadon LD (2012) The price of wind power in China during its expansion: technology adoption, learning-by-doing, economies of scale, and manufacturing localization. Energy Econ 34: 772-785

Richels RG, Blanford GJ (2008) The value of technological advance in decarbonizing the U.S. economy. Energy Econ 30: 2930-2946

Santen NR (2012) Technology investments under uncertainty: a new modeling framework for the electric power sector. Dissertation, Massachusetts Institute of Technology

US Department of Energy (DOE) (2011) Quadrennial Technology Review 2011. http://energy.gov/sites/prod/files/QTR_report.pdf. Accessed 9 September 2013

US Department of Energy (DOE) (2012) SunShot Vision Study. February 2012. DOE/GO-102012-3037. http://www1.eere.energy.gov/solar/pdfs/47927.pdf. Accessed 9 September 2013

Webster M (2002) The curious role of 'learning' in climate policy: should we wait for more data? Energy J 23: 97-119

Webster MD, Santen NR, Parpas P (2012) An approximate dynamic programming framework for modeling global climate policy under decision-dependent uncertainty. Computational Manag Sci 9: 339-362

Belfer Center for Science and International Affairs
Harvard Kennedy School
79 JFK Street
Cambridge, MA 02138
Fax: (617) 495-8963
Email: belfer_center@harvard.edu
Website: http://belfercenter.org