RC24609 (W0807-119) July 18, 2008
Computer Science

IBM Research Report

Strategic Planning for Electric Utilities under CO_2 Emissions Regulations

Stuart Siegel, Jayant Kalagnanam,
Anshul Gupta, Dharmashankar Subramanian
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Strategic Planning for Electric Utilities Under CO_2 Emissions Regulations

Stuart Siegel, Jayant Kalagnanam, Anshul Gupta, Dharmashankar Subramanian
{stus,jayant,anshul,dharmash}@us.ibm.com
Business Analytics and Mathematical Sciences
IBM T.J. Watson Research Center
Yorktown Heights, NY 10598

Executive Summary

The expected greenhouse gas (GHG) emissions regulations in the United States are likely to transform the electricity generation business in the coming decades. While preparing for a carbon constrained future, electric utilities face a far more complex capacity planning task than they are used to. This complexity stems from uncertainty in the nature of future carbon regulations as well the cost of capital-intensive and/or market-based compliance strategies.

Independent researchers have pointed out that the current analytical tools employed by utilities are inadequate for formulating an optimal long-term capacity and compliance planning portfolio under such uncertainty. Additionally, current methods do not allow the explicit modeling of tradeoffs between the usually conflicting goals of minimizing expected portfolio costs and minimizing the uncertainty in portfolio costs. This poses a significant risk to electric utilities and their stakeholders since long term planning is a critical aspect of the utility business due to the long economic lifetime (30-60 years), long development and lead-time, and large capital costs of the generation and transmission infrastructure. The plans that are being drafted today must consider the impact of GHG emissions regulations because optimal portfolios of generation capacity with and without such considerations are likely to be significantly different from each other.

In this report, we describe a novel patent-pending integrated capacity planning framework to aid utilities doing long term capacity planning under the uncertainties of a carbon-constrained future. Our framework has the capability of modeling uncertainties in the inputs to yield a portfolio with the least expected costs within the risk tolerance of a decision maker. Unlike most current planning approaches that rely on the analysis of a relatively small number of scenarios or portfolios, our tool is aimed at generating optimal risk-modulated portfolios based on optimal abatement investments and their timings. We believe that electric utilities would need this capability for effectively planning their long-term capacity and environment compliance activities.

1. Introduction

There is an emerging international consensus around the impact of anthropogenic greenhouse gas (GHG) emissions on the earth's climate (IPCC Report 2007)[1] and the need for limiting such emissions. The predominant source of GHGs, mainly CO_2, in the United States is fossil-fuel based electricity generation (see Figure 1)[2].

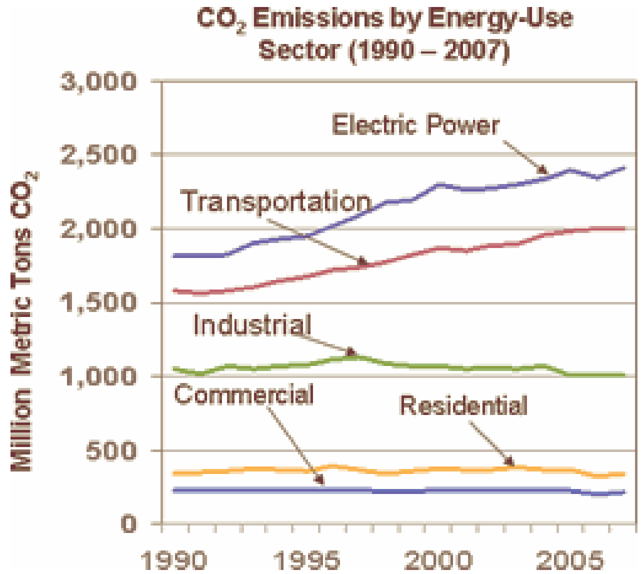

Figure 1: US CO_2 emissions by sector.

It is widely expected that in not too distant a future, the US will join most industrialized nations in enacting binding legislation to limit GHG emissions. Several state and regional initiatives in the US and Canada are already gaining momentum. The Regional Greenhouse Gas Initiative (RGGI), which is set to begin in 2009, is a commitment by at least 10 northeastern states to cap regional CO_2 emissions at 1990 levels by 2014 and to reduce them by 10 percent below that level by 2018. In 2006, California passed legislation requiring a 25-percent reduction in CO_2 emissions by 2020. The Western Climate Initiative, modeled after RGGI, has set a goal of bringing regional emissions to 15 percent below 2005 levels by 2020 by establishing a market mechanism. It is, therefore, imperative that US electric utilities equip themselves for future regulatory compliance with respect to GHG emissions.

[1] http://www.ipcc.ch/
[2] http://www.eia.doe.gov/oiaf/1605/flash/flash.html

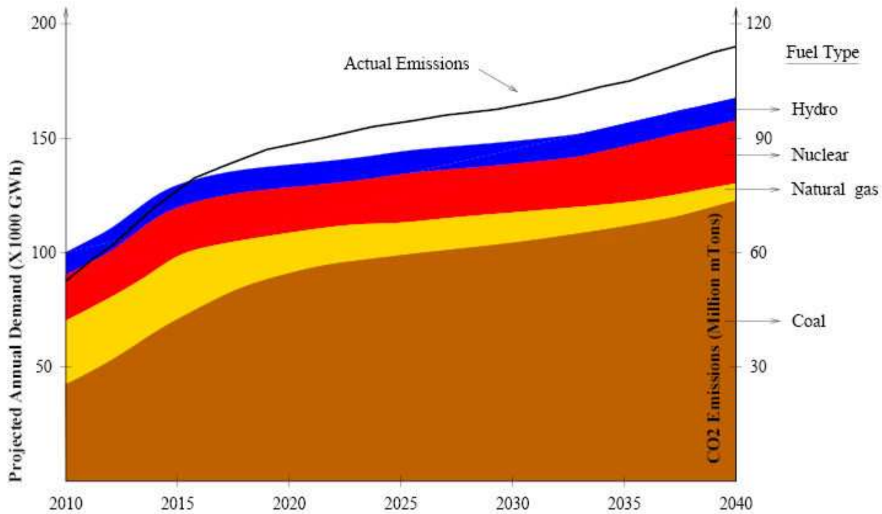

Figure 2: A 30-year CO_2 emissions and fuel mix profile of a hypothetical utility under the assumption of no GHG emissions cap.

Long term generation capacity planning is a critical aspect of the electric utility business due to the long economic lifetime (30-60 years), long development and lead-time, and large capital costs of the generation and transmission infrastructure. The plans that are being drafted today must consider the impact of GHG emissions regulations, because optimal portfolios of generation capacity with and without such considerations are likely to be significantly different from each other. Figures 2 and 3 illustrate how the long-term electricity generation portfolios of a hypothetical utility might look like in a carbon-oblivious and a carbon-constrained scenario, respectively. In the carbon-oblivious scenario of Figure 2, increasing demand is met by an increasing dependence on coal. As a result, while demand grows a little over 50% in 30 years, CO_2 emissions nearly double over the same period.

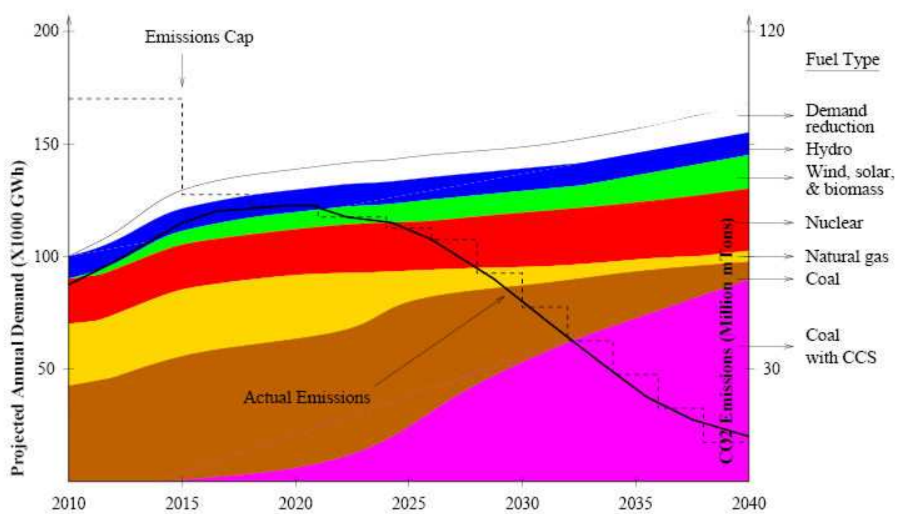

Figure 3: A 30-year CO_2 emissions and fuel mix profile with a long-term resource planning approach that considers impending GHG emissions caps.

The electricity generation portfolio of the same utility might look very different, like the one shown in Figure 3, under a hypothetical GHG regulatory regime that imposes a cap on emissions (shown as the dotted line in the figure). The generation mix shown in Figure 3 must be carefully optimized to minimize the capital and operating costs, while meeting constraints on demand satisfaction and caps on emissions. It will most certainly need to rely on a meticulous strategy of choosing various abatement options, fuel mixes, and capital investments in new capacity and retrofits, along with optimal timing of these decisions. In reality, the problem of constructing an optimal portfolio is even more complex because of uncertainties surrounding the regulations and the various costs involved.

Until recently, utilities would plan future power plant constructions based on their projections for relatively few variables with a moderate degree of unpredictability, such as fuel availability and costs, demand projections, and environmental regulations, etc. Lately, however, they are rapidly finding themselves in an extremely complex and uncertain planning environment. While there is a consensus on the inevitability of binding federal regulations to limit GHG emissions, there is still considerable discussion and speculation regarding the nature and timing of such regulations, their impact on electric utilities, and the feasibility and costs of various abatement options. There are three key new sources of uncertainty that the utilities must consider: (a) direct impact of regulations, in terms of timing and costs of GHG emissions regulations; (b) indirect impact of regulations, in terms of changes in fuel prices, fuel transportation costs, wholesale electricity market prices, air pollutant permit prices, and demand; and, (c) future costs, performance, and adoption rates of new technologies for emissions abatement (such as, carbon capture and storage or CCS), generation (such as, solar PV), and consumption (such as, plug-in hybrids, and demand-response enabled with advanced metering infrastructure). A decision maker at a utility faced with the task of developing a strategic resource plan that is cognizant of pending carbon regulation needs to be able to develop a portfolio of different generation units over the planning horizon (typically 25-30 years) that minimizes expected cost (or other risk measures) while complying with CO_2 limits. To develop such a portfolio, the decision maker needs to design an investment policy that considers the significant regulatory, financial, political (e.g., will nuclear ever by socially palatable again in the U.S.) and technological uncertainties.

In this report, we discuss a novel patent-pending stochastic optimization based decision support system that allows a decision maker to model this problem and identify a portfolio of abatement options that minimizes the cost of abatement over a long time horizon. We provide a comparison of a typical scenario based analysis that many utilities currently use, and contrast this against an optimal resource plan generated using our integrated analysis framework. Our tool can consider various abatement options such as (a) shifting the fuel source mix away from coal, (b) investing in non-fossil sources like nuclear or renewables, (c) investment in ICCG (integrated gasification combined cycle) plants with CCS (carbon capture and storage), (d) investment in demand-side management programs including smart metering infrastructure, (e) enabling and encouraging increased profusion of distributed generation using wind and solar, (f) CCS retrofit for existing fossil fuel units, (g) entering land-use contracts for afforestation and conservation tillage to offset CO_2 emissions, (h) strategic buying, selling, and banking of CO_2 allowance over time, (i) the extent of power purchased, vis-à-vis, power generated in-house, in order to satisfy demand, etc. The system is designed to allow for easy addition of other abatement options for which sufficient cost and performance models become available. It is capable of modeling the aforementioned regulatory, market, and technological uncertainties. A key feature of our system is that it allows a decision maker to supply an explicit risk-based[3] objective function to the tool to

---

[3] We consider symmetric risk measures (e.g. standard deviation as a measure of variability), as well as asymmetric, tail-risk measures that are mature in the financial literature to model worst-case outcomes, such as Value-at-Risk (VaR) and Conditional-Value-at-Risk (CVaR).

optimize. In other words, it would create a portfolio with minimum expected cost such that the risk (probability) of the total cost exceeding its expected value, or any user-specified value, by a certain given factor is within a user specified tolerance.

2. Capacity Planning Challenge in a Carbon Constrained Future

The most common resource planning methods used today are based on scenario analysis over a set of hand-crafted portfolios. Uncertainty is handled by simulations (Monte-Carlo techniques) over a small number of predefined scenarios in order to generate the probability distributions around cost and compliance estimates. Portfolios are then ranked across the above, limited number of predefined scenarios using measures such as expected cost or other risk measures from the finance literature such as, Percentiles, Value-at-Risk, or Conditional-Value-at-Risk.

Consider a utility with a current generation mix close to that at the beginning of the 30-year planning horizon shown in Figures 2 and 3. The utility faces a number of difficult choices over the next 30 years. Its optimal behavior is highly dependent on the aforementioned uncertain factors. For the sake of presentation tractability, let's explore just three of these uncertainties:

1. Will there be a CO_2 emissions cap placed on the utility and if so, what will the price of an emissions allowance be?
2. What will be the feasibility and cost of building new coal units, as well as retrofitting existing coal units with carbon capture and storage (CCS) equipment.
3. Will nuclear base-load generating capacity plants be a viable investment option?

For the CO_2 regulatory uncertainty, we'll consider three possible scenarios: (i) no emissions cap, (ii) a "moderate" emissions cap with moderately priced allowances, and (iii) a "severe" emissions cap with costly allowances. For the CCS uncertainty, we'll consider two possible outcomes: (i) CCS is not viable, (ii) it's viable at "moderate" cost. For the nuclear availability question, we assume two possible outcomes: (i) nuclear is viable (politically, economically, or otherwise) and (ii) it is not.

We can view the decision problem facing the utility using a decision tree tracing each of the possible realizations of the future and the likely preferred portfolio strategy. This is shown in Figure 4. Given the projected demand growth, the utility will need to invest in new generating capacity within the planning horizon. The choices (nuclear, NGCC, PC, IGCC with CCS, renewables, demand-side management, etc.) will depend on how the "future plays out" with respect to the uncertain inputs. Figure 4 shows how the optimal strategy differs depending on the path taken in the tree. Each path represents one realization of our artificially bounded future. Take, for example, the interesting case of "severe" CO_2 regulations. In this case, there might be a real strategy-based argument for early investment in zero or relatively low CO_2-emitting generation technologies and early retirement of older, less efficient units. Such a strategy would allow the utility to obtain revenue by selling valuable CO_2 allowances in the market. In a "moderate" regulatory scenario, on the other hand, the utility might get away by simply supplementing an almost carbon-oblivious portfolio with demand reduction, modifying fuel mix, some investment in renewables, and purchasing allowances.

Utility planners often use this type of gating analysis to come up with a narrow set of possible portfolios based on their stated expectations of certain key parameters. Then the selected portfolios are subjected to more detailed analysis to establish bounds on costs and other values of interest to regulators and ratepayers.

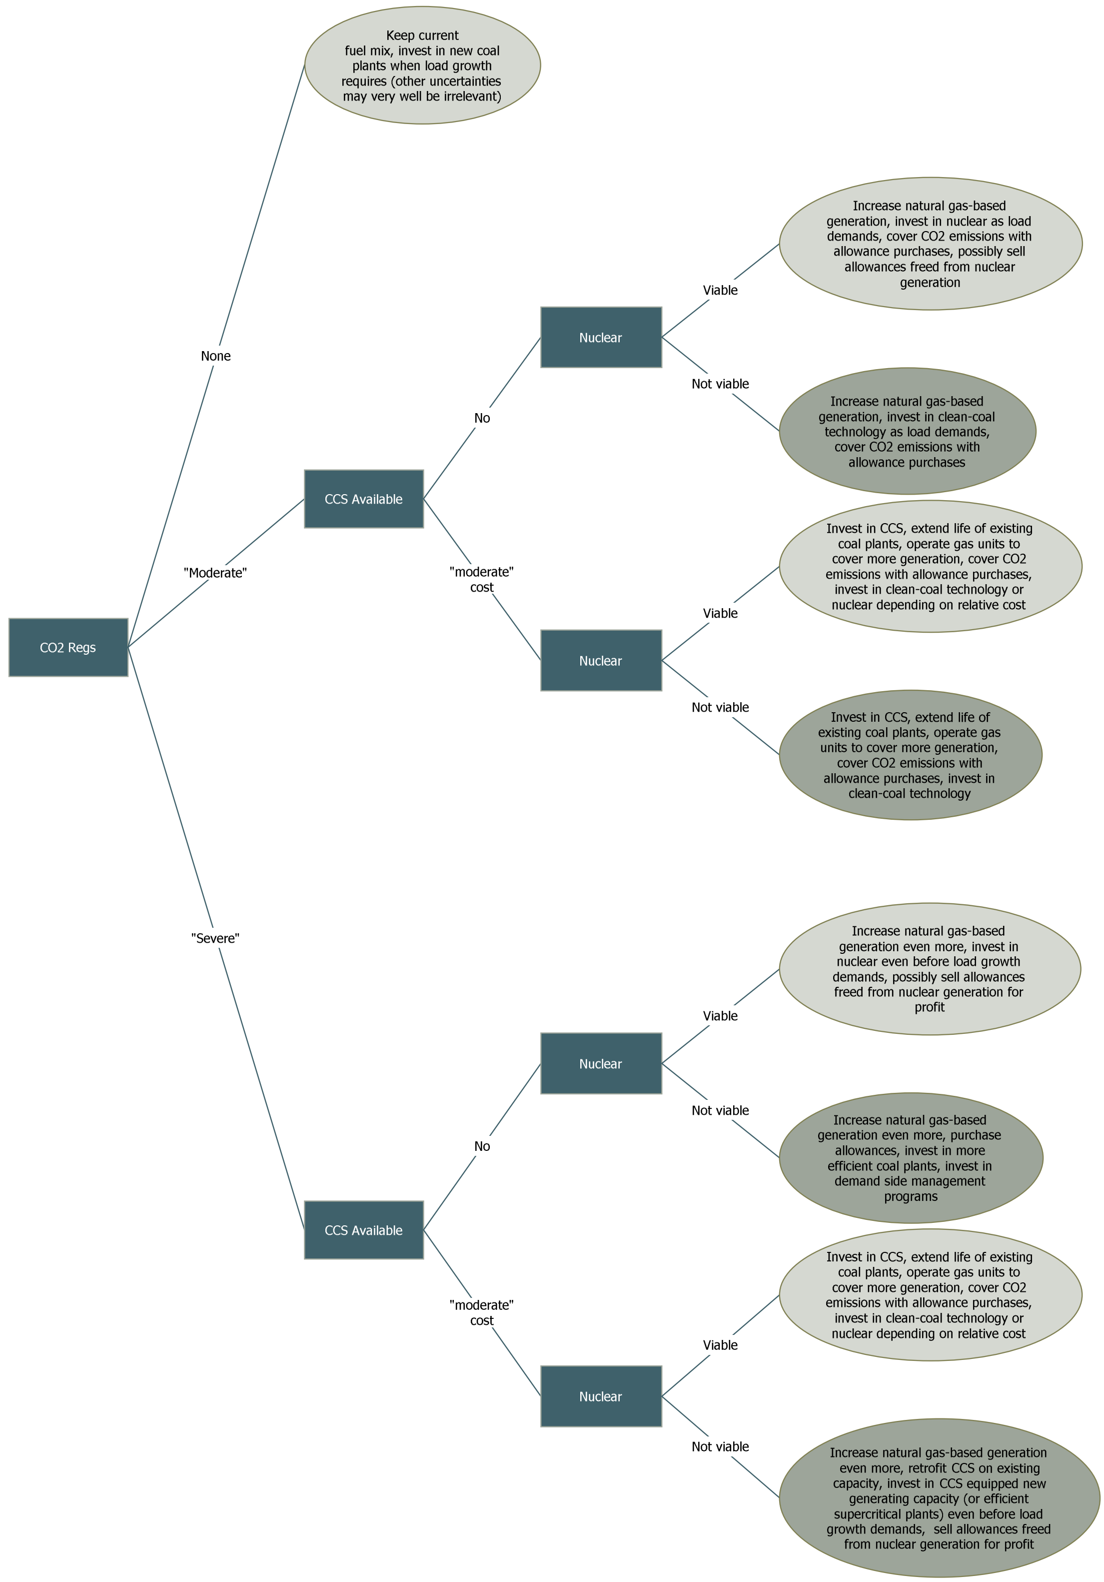

Figure 4: Decision tree showing possible preferred outcomes under input parameter uncertainty for three parameters.

There are a number of limitations with looking at the decision problem this way. First and perhaps most obvious one is that, in the interest of problem tractability, the number of uncertain inputs in the model has to be fairly limited. While such an approach for capacity planning might have been practical until now, the increase in the number of uncertain inputs due to impending emissions regulations makes the number of scenarios intractable. Although computers can assist the utility decision maker enumerate the possibilities, hand-crafting the portfolios under each scenario, which is a common practice today, would not be feasible. Consider the case of just five uncertain inputs with five possible outcomes each. In this case, an analyst would need to create portfolios for over three thousand scenarios. Another problem with this approach is that there is no reason to expect that the odds of a particular outcome for an uncertain input are independent of the values of other uncertain inputs. For example, it's plausible that if "severe" CO_2 regulations did come to pass, political and popular support for nuclear plants might change since they are based on a mature zero-emitting technology and can displace existing large coal plants with minimal grid disruptions. Another limitation of the approach is that the important time dimension of the problem tends to get lost. For all of the aforementioned uncertainties, capturing the "when" is as important as modeling the "if", "how much", and "what" aspects. The reason is that the relative likelihood of various outcomes is strongly related to time. For example, the probability of CCS being a viable option in 2010 is relatively low, but significantly higher in 2020. The consequence of these limitations is that a decision made based on such an approach may not be robust to all the possible outcomes that may occur over time. Since only the analysis and selection of portfolios is automated, not their creation, a preferred portfolio chosen by this approach may be highly suboptimal with respect to both the expected cost and the cost uncertainty.

3. An Integrated Capacity Planning Approach

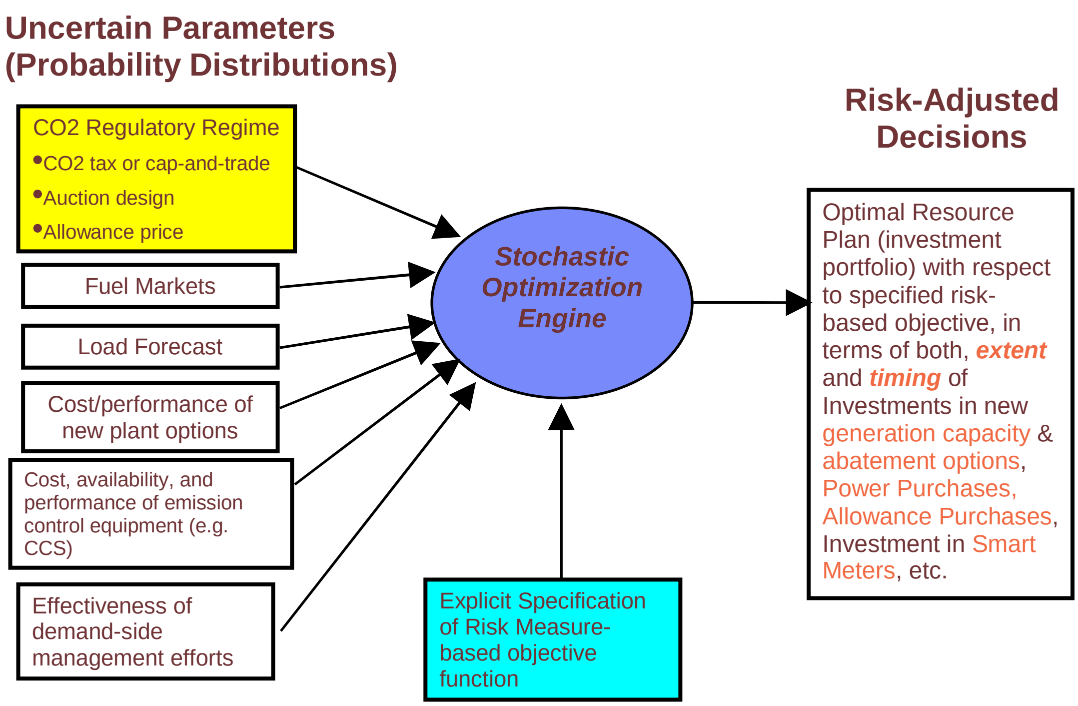

Figure 5: An overview of the integrated analysis framework.

In order to address some of the aforementioned limitations, we have constructed a decision support tool that can handle practically any number of uncertain inputs in a continuous, integrated, and time-indexed fashion without becoming computationally intractable. Our tool allows for the explicit modeling of correlated uncertain inputs, using probability distributions and dependence structures like correlations, and uses formalized optimization methods such as mixed-integer linear programming and stochastic optimization to provide a utility decision maker with near-globally-optimal plant operations and capacity investment portfolios that are robust in the face of all modeled uncertainties. A decision maker at an electric utility could use this tool to meaningfully translate information about variation in portfolio costs across multiple scenarios into the selection of a single preferred portfolio. Furthermore, the tradeoff between the twin objectives of minimizing expected portfolio costs and minimizing uncertainty in portfolio costs (in terms of financially meaningful risk measures of the cost distribution) is explicitly handled in this framework. Figure 5 provides an overview of the framework.

At the core of this framework is a stochastic optimization engine that provides the ability to optimally pick the portfolio that minimizes the expected cost of compliance for a given utility over a set of uncertain compliance options. In the framework that we have developed, uncertainties are modeled using probability density functions that may or may not vary with time. For example, Figure 6 shows how the CO_2 regulatory uncertainty is modeled as time-dependent (hypothetical, selected for purely illustrative purposes) distributions. The first curve marked "2010-2012" shows that the modeler expects that in the near term it is most likely that CO_2 emissions will be priced at zero (i.e., no carbon cap or tax at all) with some small probability that CO_2 could be priced up to $3/ton emitted. In contrast, the second curve marked "2016-2018" represents a much different expectation that, by that time, CO_2 regulation will come to bear and the allowance price or tax will be in the $15-25/ton range with an expected value of $20/ton. Similar uncertainty distributions could be specified for more time periods and other uncertain parameters.

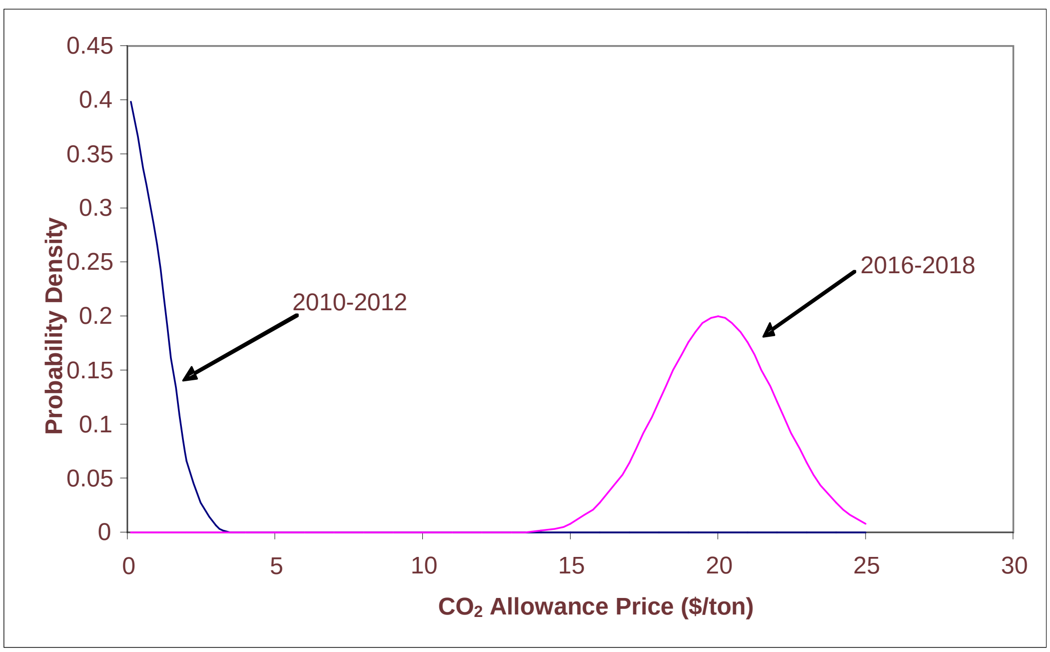

Figure 6: Example of modeling CO_2 regulatory uncertainty using probability density functions.

It is important to point out that simply providing the capability of modeling uncertainty this way is necessary but not sufficient for enabling better decision making. The specification of the forms of the distribution is potentially a difficult job and would presumably need to be done with careful consultation between utility planners, forecasters, and regulatory overseers to make sure that reasonably realistic distributions are selected. This process in itself is potentially a beneficial exercise for utilities because independent and careful scrutiny of inputs before the analysis makes the ultimate modeling results based on said inputs more defensible.

Assuming that all uncertain parameters (load growth, technological costs, fuel costs, regulatory uncertainties, etc.) are modeled in a similar fashion as the CO_2 allowance price example, our decision making framework can accept these inputs and pass them through a stochastic optimizer engine. The engine will select a least-cost (expressed in present value terms) time-dependent portfolio for meeting future load and CO_2 regulatory requirements. For example, for the hypothetical utility represented in Figure 3, the optimization engine will yield a portfolio with the least expected cost that can be operated to meet the expected demand and emissions cap at any given time during the planning horizon using the available fuel options. The options do not have to come bundled together as a given pre-selected portfolio. The optimization engine that we have developed treats each option (CCS, natural gas re-dispatch, investment in renewable sources, etc.) as discrete, selecting each automatically to form an optimal portfolio of investment and plant operations decisions. The solution consists of the "best" mix of the options *and* the optimal timing to introduce each of the options. Thus, it is ideally suited for addressing strategic questions, such as, "is it better to invest in emissions lowering technologies now or later?"

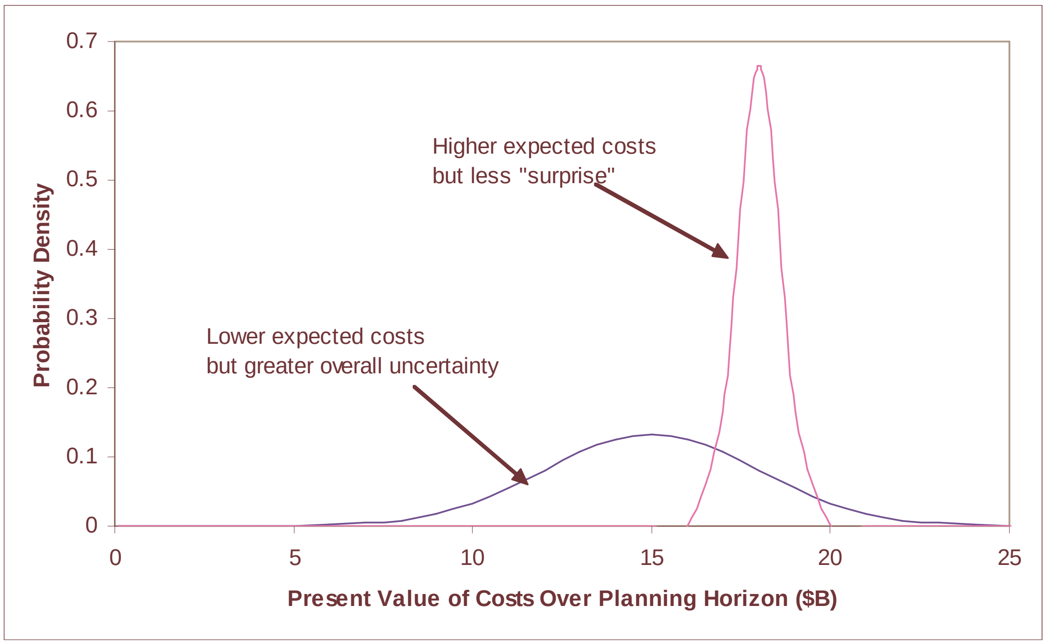

Figure 7: Distribution of costs over planning horizon for two optimal (depending on desired risk vs. expected value tradeoff) investment portfolios.

We now discuss the second key feature of our analysis tool. Just like the uncertain inputs are represented by probability distributions, the cost of a portfolios generated by the tool itself is a probability distribution and is a part of the tool's output. Figure 7 shows the distribution of the present value of the capital and operating costs associated with two hypothetical long-term electricity generation capacity portfolios. The tool could generate either of the portfolios as the "optimal" portfolio, depending on the form of the objective function specified in the optimization model. The first, with an expected value of about $15B, represents the case where the expected cost is low, but there is a higher risk of the cost substantially exceeding the expected cost. The second distribution, with a mean of $18B, has a higher expected cost, but a much smaller variance than the first. This added certainty that costs will not vary substantially from the expected value may be valuable enough to many utilities and regulators to make the second portfolio the preferred choice even though it has a higher expected cost. Our tool allows the decision maker to explicitly specify this desired level of tradeoff between expected costs and uncertainty in costs and the optimization engine will factor this in when solving for the optimal portfolio. In addition to symmetric risk measures like standard deviation that measure variability ("surprise"), our tool also allows the decision-maker to focus on asymmetric, tail-based risk measures such as Value-at-Risk, or Conditional-Value-at-Risk, if the objective is to minimize (or contain, within some user-specified tolerance) the worst-case value of the cost distribution, say the 99^{th} percentile of the cost distribution.

It should be noted that the optimization model can be used to provide two types of outputs. Firstly, for a given scenario, we can augment the current practice of using Monte Carlo simulation with risk-aware, stochastic optimization to generate the *optimal timing* of investments to manage risk at minimum expected cost. Secondly, across all possible scenarios, we can generate an optimal portfolio, both in terms of the *optimal mix (extent)*, as well as, *optimal timing* of investments, to manage the risks associated with carbon compliance at minimum expected cost.

4. Summary and Conclusion

The widely expected impending greenhouse gas emissions regulations in the United States are likely to fundamentally transform the electricity generation business in the coming decades. While preparing for a carbon constrained future, electric utilities face a far more complex generation capacity planning task than they are used to. This planning complexity stems from the highly uncertain primary and secondary effects of the future emissions regulations on virtually all aspects of the electric utility business. Current analytical tools are inadequate for both portfolio selection under a large number of uncertain inputs and for explicitly modeling tradeoffs between the usually conflicting goals of minimizing expected portfolio costs and minimizing the uncertainty in portfolio costs. In this report, we have described a novel patent-pending integrated capacity planning framework to aid utilities plan their generating capacity under the uncertainties of a carbon-constrained future. Our framework has the capability of modeling uncertainties in the inputs to yield a portfolio with the least expected costs within the risk tolerance of a decision maker. We believe that electric utilities would need this capability for effectively planning their long-term generating capacity.