Energy Efficiency Portfolio Risk Management: A Systematic, Data-Driven Approach for Timely Interventions to Maximize Results

Richard Ridge, Ridge & Associates, Alameda, CA
Steve Kromer, Teton Energy Partners, Oakland, CA
Steven Meyers, Rational Energy Network, Austin, TX
Doug Mahone, Heschong Mahone Group, Inc., Fair Oaks, CA
Jay Luo, Pacific Gas and Electric Company, San Francisco, CA
Valerie Richardson, Pacific Gas and Electric Company, San Francisco, CA
Rafael Friedmann, Pacific Gas and Electric Company, San Francisco, CA
Nick Hall, TecMarket Works, Oregon, WI

ABSTRACT

A systematic, data-driven approach is being used to maximize the net benefits of Pacific Gas and Electric Company's portfolio of energy efficiency programs. The analysis centers on identifying and minimizing riskier elements affecting portfolio performance and allowing informed optimization of resources. Given the recent tripling in funding and the stretch energy saving goals established for the 2006-2008 period, an approach providing actionable results for timely corrections to the portfolio implementation is needed.

This paper describes an approach for identifying, quantifying, and managing risk that maximizes the probability that PG&E will meet its energy and demand savings targets. The keystone of our systematic risk management approach is Monte Carlo simulation using Crystal Ball (CB) software. CB automatically calculates thousands of different "what if" cases, saving the inputs and results of each calculation as individual scenarios. Analyzing these scenarios reveals the range of possible outcomes, their probability of occurring, which inputs have the most effect on your model results and where program managers should focus their activities to manage risk.

This paper provides concrete examples of how both primary and secondary data have been used to modify various parameters (e.g. net-to-gross ratios, installed units) that adjust the forecast and assess the probability of not achieving the energy and demand targets. The simulation results guide the selection of actionable strategies to manage this risk. Continuous monitoring of key portfolio performance indicators (e.g. installation counts and levelized cost) help determine the success of each particular strategy.

Background

Probabilistic Monte Carlo simulations have historically helped analyze and value complex energy systems such as demand response resources (Violette and Freeman 2004), grid portfolio risks (Schilmoeller 2004), and private (non-utility) energy-efficiency portfolios (Kromer, Sezgen, and Meyers 2004). During 2006, the TecMarket Works Team (Team)[1] was under contract to the CPUC to develop evaluation protocols (Hall et. al 2006). As a part of this effort, a subset of the TecMarket Works Team (Sub-Team) was asked to design a system for recommending to the California Public Utilities Commission – Energy Division (CPUC) how to allocate resources for ex post evaluations across the

[1] Hall, Nick, Johna Roth, Carmen Best, Sharyn Barata, Pete Jacobs, Ken Keating, Ph.D., Steve Kromer, Lori Megdal, Ph.D., Jane Peters, Ph.D., Richard Ridge, Ph.D., Francis Trottier, and Ed Vine, Ph.D.

portfolio's energy efficiency programs. This report (Hall, Jacobs, and Kromer 2006) also informed the evaluation planning process regarding which programs should be more or less rigorously evaluated. If a specific program delivery and measure mix strategy had already been rigorously evaluated then allocating new dollars to repeat these studies may not be the best use of evaluation resources.

In typical evaluation planning approaches, expert evaluation professionals assess results from past studies to prioritize the programs delivery strategies and measure combinations. In these efforts, the focus of the review is on previous evaluations and the level of accuracy and precision (reliability) of the projected savings. If there are gaps in information that can be obtained with new efforts, then those efforts are prioritized and evaluation funds are allocated to fill those information gaps. The overall goal of these efforts was to conduct evaluations which provide a more accurate estimate of savings for each evaluated program and the portfolio of programs as a whole. These planning efforts are not new to professional evaluators and experienced evaluators are accustomed to and comfortable with these approaches. However, with a potential 250 programs encompassing hundreds of different measure offerings and delivery approaches, developing an expert opinion regarding the evaluation planning approach becomes more problematic. Keeping track of the gaps in evaluation information at the measure level, as influenced by the numerous program delivery strategies within which the savings are achieved suggested that a new, more formalized and potentially more accurate evaluation planning approach was needed.

The Sub-Team decided that a Monte Carlo simulation of the CPUC's portfolio could systematically track and quantify the hundreds of relevant data gaps. These results would then guide the allocation of evaluation resources cost-effectively to the most deserving elements of the portfolio. The overall objective of the approach was to identify the programs and program measures that needed the greatest attention so that evaluation study results would more reliably document the achieved savings.

The scope of the risk analysis effort was limited to identifying sources of uncertainties in the ex-ante savings estimates used by the IOUs during program design in order to inform the process of prioritizing and allocating evaluation resources. The project was not designed to estimate the magnitude of the uncertainty in the overall portfolio savings, or assess the probability that the program implementers would meet their energy savings goals (that effort came later and is the primary focus of this paper). Also, the issue of the amount of uncertainty reduction resulting from a particular evaluation approach and the costs of achieving this level of reduction were not addressed. The results of the CPUC-sponsored risk analysis efforts were used first to identify major sources of uncertainty in portfolio savings estimates and then to allocate evaluation resources so that the program evaluations that had the greatest potential to reduce estimation error for the programs with the greatest risks received considerably more evaluation resources.

To assess the level of risks, the Sub-Team set uncertainty metrics for three conditions associated with the program-specific measure-specific savings estimates, with individual metrics set for kWh, kW and therm savings. The uncertainty metrics used in this assessment include:

1. The uncertainty around the measure-specific energy savings (unit energy savings) projections, for kWh, kW and therms,
2. The uncertainty around the program's estimated installation counts as recorded in the program tracking system by each delivery strategy,
3. The uncertainty around the program-projected net-to-gross ratios for each delivery strategy.

For each measure and for each program, the uncertainty boundaries (high boundary and low boundary) associated with these three metrics were estimated for each of the types of energy to be saved for each program and program measure group. The process of estimating the high and low boundaries for these metrics for each program and for each measure group involved the use of past evaluation findings combined with three rounds of professional opinion adjustments, with changes to the estimated boundaries made after each round of the reviews.

Once the uncertainty boundary estimates were established during this multi-step review process, the metrics and their associated boundaries were distributed across a standard distribution curve. Because of the complexity of this comparison process the Sub-Team used risk analysis software to handle the several hundred thousand comparative calculations needed for scoring the relative uncertainties that would guide the evaluation efforts. The software used is Decisioneering's *Crystal Ball®* (CB) version 7 (www.decisioneering.com). The results of these calculations were then used to identify the program and measure groupings that presented the greatest levels of uncertainty for the portfolio as a whole and within each program. This information was then used to establish rigor levels and allocate resources to studies that would best help reduce the level of energy savings uncertainty from the evaluation efforts, and therefore provide a more reliable estimate of the savings from the portfolio. As a result of this effort, less evaluation resources and lower levels of evaluation rigor are planned for programs that do not add substantial risk to the energy savings estimates, and higher rigor, higher cost studies are focused on the programs and measures that contribute the highest level of risk.

The Utility Perspective

Although the CPUC used a Monte Carlo simulation of the state's energy-efficiency portfolio to allocate resources for *ex post evaluations*, the Pacific Gas and Electric Company (PG&E) realized that, with one important change in perspective, this approach could help manage their energy efficiency portfolio. Recall that the Sub-Team, from an ex post evaluation perspective, was concerned about the uncertainty around the program's installation counts as *recorded* in the program tracking system by each delivery strategy. However, from the perspective of managing portfolio risk, PG&E was more concerned about whether it could meet its *forecasted* installations over the three-year period (2006-2008). That is, "recorded installation counts" was replaced with "forecasted installation counts" with a great deal more uncertainty around the latter. Thus, for a given PG&E program, the basic planning equation for achieving its goals is:

Gross Unit Energy Savings x *Forecasted Number of Measures Installed* x *NTGR*

These three parameters became the central focus of PG&E's effort to manage its portfolio.

There are a variety of sources of uncertainty and threats to each of these parameters. Some of the key elements of uncertainty that contribute to risk are:

- That the program theory is incorrect and cannot achieve the forecasted unit installations
- That the program is not faithfully implemented leading to a failure to achieve expected unit energy savings, unit installations or NTGRs
- That PG&E's original, ex ante, estimates of unit energy savings, forecasted unit installations, and NTGRs were flawed
- That CPUC impact evaluators might develop flawed estimates of unit energy savings and NTGRs

PG&E's Customer Energy Efficiency (CEE) Department is in a position to obtain a great deal of information that can be used to identify risk and develop actionable strategies to minimize the risk of not achieving portfolio energy and demand goals. A variety of analyses can be conducted to provide information that can be used to change strategies and program designs and even reallocate program dollars in a way that maximizes PG&E's chances of meeting the portfolio savings target. The general framework for managing program and portfolio risk is presented in Figure 1. Each of these elements of uncertainty is briefly discussed below.

The first area of risk is that the savings estimated by the impact evaluation team, under contract to the CPUC, might be either over- *or* under-estimated. One way to manage this risk is to measure the right things in the right way. That is, PG&E must make sure that their programs are adequately described so that all savings, including any spillover, can be targeted by the impact evaluation team and that the attribution of these savings to PG&E programs is correctly handled. PG&E must be convinced that the methods recommended by the impact evaluation team are capable of producing reliable estimates of gross and net savings impacts. Consider a case in which an evaluator suggested using a billing regression analysis as a way of estimating savings for large industrial customers and asking a single question upon which to base a net-to-gross ratio (NTGR). In such a case, PG&E should argue strongly against this approach to estimating gross impacts because of the low signal to noise ratio and the approach to estimating the NTGR because of unreliability. The discussion regarding the best approach will be conducted during the research planning phase for each impact evaluation in close collaboration with PG&E staff.

The second area of risk is that the theory underlying a given program is not supported in the social science literature. For example, a program could, among other things, provide one hour of sales training for participating retailers and consumer education, including in-store point-of-sale displays. If a review of the evaluation literature found that one hour of sales staff training was not sufficient and that expected sales staff turnover of 30 percent per year would require a continuing investment that was not originally anticipated by the program designer. In such a situation, one could recommend more staff training along with the necessary budget to provide on-going training of new sales staff.

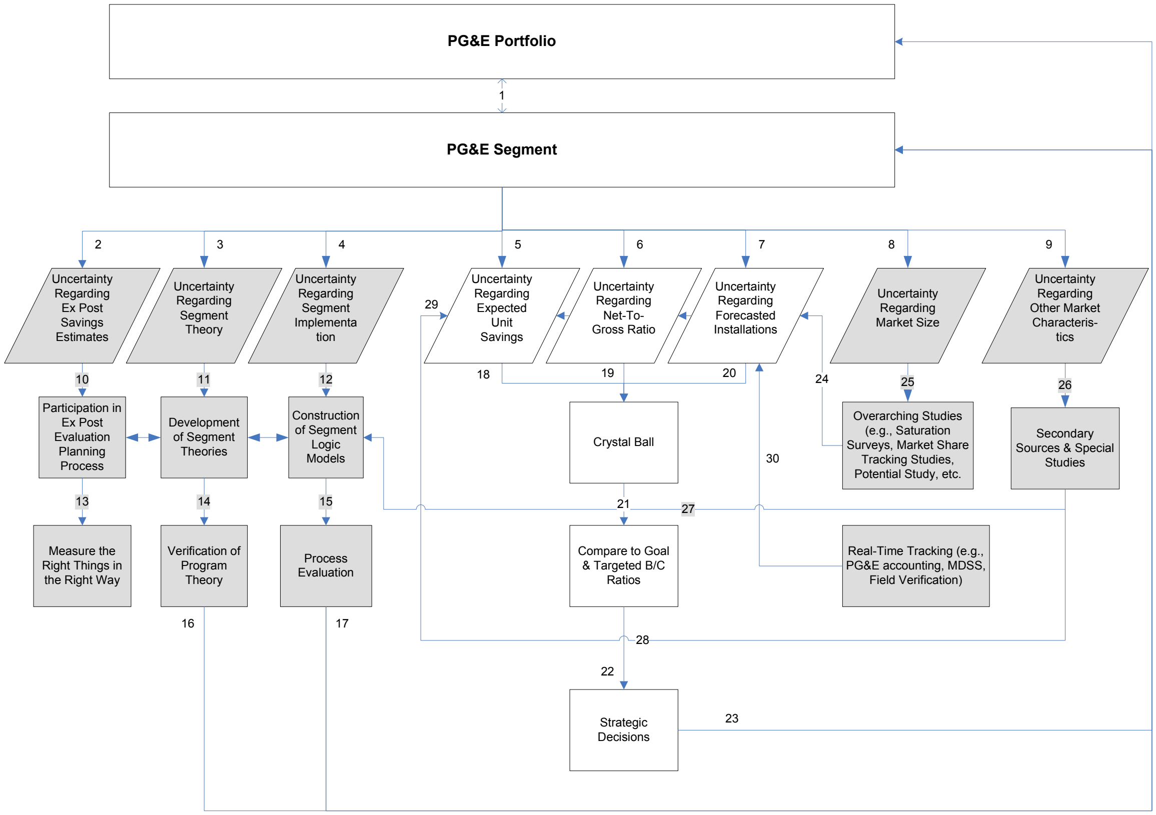

Figure 1. The PG&E Portfolio Risk Management Model

One way to mitigate this risk at the beginning is to develop a program theory and logic model. PG&E has already developed program theories and logic models for each of its segments. These program theories and logic models can: 1) provide a sanity check on underlying theories, 2) support process evaluations designed to support continual improvement in program design and implementation, 3) be used as a starting point for impact evaluations that estimate energy and demand impacts, 4) be used to explain segment-level efforts to others. The results of developing a program theory and logic model can be used to make mid-course corrections or even modify the design of the program.

The third area of uncertainty and risk is that the program as designed is not faithfully implemented in the field leading to less than expected savings. For example, if a program, among other things, was supposed to provide two hours of classroom training to retail appliance store sales staff but actually provided only 20 minutes of in-aisle training. This would make the sales staff less able to inform customers about energy efficient products resulting in fewer purchases of Energy Star appliances. In such a case, the evaluation team would recommend that each training session consume the full two hours. The best approach to mitigating this risk is to conduct a process evaluation, the primary objective of which is to help program designers and managers structure their programs to achieve cost-effective savings while maintaining high levels of customer satisfaction.

The next three parameters are directly tied to the portfolio energy and demand goals filed by PG&E: 1) ex ante unit savings, 2) ex ante net-to-gross ratios, and 3) forecasted installations. It is possible that one of more of these parameters is biased. In addition, whether biased or not, there is some uncertainty surrounding each. Each of these three parameters can be refined using: 1) regular inputs from the impact evaluation about gross impacts and NTGRs, 2) any changes in the program design resulting from an examination of the program theory, 3) any mid-course corrections in the delivery of the program recommended by the process evaluation, 4) more accurate information about the size of the market and the key market actors and how they interact to influence the purchase and installation of energy efficient equipment, and 5) regular inputs from the program-tracking database. We've already briefly discussed the first three. The fourth, better information about the market can be acquired by conducting special studies of cost-effective energy efficiency potential, saturation surveys, and market characterization studies. Finally, monitoring of the program-tracking database can help to determine if PG&E is on track with its various forecasts.

One could take all this information and, for example, through a Delphi technique, agree upon any necessary changes in the program design and delivery, a re-allocation of dollars to other programs or segments, and changes ex ante assumption regarding unit savings, NTGRs, or forecasted installations. There is, however, another more systematic approach using Crystal Ball (CB), software that provides an easy way to perform Monte Carlo simulations in Excel spreadsheets. CB automatically calculates thousands of different "what if" cases, saving the inputs and results of each calculation as individual scenarios. Analyzing these scenarios reveals the range of possible outcomes, their probability of occurring, which input has the most effect on the model and where program managers should focus activities to manage risk. CB can help portfolio managers maximize their chances of achieving their savings goals. . For example, a given program and measure may estimate the NTGR to be 0.7. However, if early results from the impact evaluation suggest that the true point estimate is 0.60, plus/minus 0.1, then one could re-run the simulation assuming the full range of NTGR from 0.5 to 0.7 and quantify the impact on the program as well on the entire portfolio. If the impact was significant, then one might want to consider a change in the design or delivery of the program in an effort to increase the NTGR. Or, one could decide to shift funds to other measures and segments with similar gross impacts but with higher NTGRs. Or, perhaps a potential study suggests that the potential for savings for a given measure in a particular market is only half what was originally forecasted. Again, one could re-run CB to quantify the impact on the program and the entire portfolio. If the impact was significant, then one might consider shifting funds to areas with greater potential. On-going recalculation of program and portfolio savings with CB can provide on-going feedback about key parameters to which program planners can respond in various ways to keep the portfolio on track to achieve its energy and demand goals.

The PG&E Case

PG&E, committed to identifying and quantifying risk in its 2006-2008 portfolio, retained a team of consultants (the Risk Team), using the basic framework outlined in above, to adapt the CPUC model to PG&E's specific needs. The initial scope of work for the Risk Team included:

- Reconstructing the CPUC portfolio risk using PG&E's updated program definitions
- Updating the model inputs to reflect specific PG&E conditions
- Updating the model to match PG&E's unique "segment" definitions
- Allocating Local Government Programs (LGP) and 3rd Party Programs to the appropriate segments
- Creating a user-interface "dashboard" for easy review of risk data
- Collecting additional information from program mangers and EE workers (audit teams)
- Further updating the model and adjusted certain segments to match segment potential using work papers and other market impact assumptions.
- Running several additional simulations to incorporate newly available data into the results.

- Reporting results throughout PG&E and continued to update the model as new data became available

As noted above, the Sub-Team created a California-wide portfolio risk model for use in assisting the allocation of evaluation resources. That model was constructed based on the specific requirements of the CPUC evaluation activity. Among the constraints that drove the initial model design were, 1) the model would be based on the current version of the Avoided Cost Workbooks that the CPUC required for IOU portfolio filings (ACW v2d), 2) uncertainty estimates would be applied uniformly across all IOUs, varying only by general categories of measure group type and delivery strategy, not by IOU programs, 3) the overall model would be focused primarily on kWh, kW and therms, and 4) no program cost information would be included.

At least part of the justification for applying portfolio uncertainty modeling inside PG&E, and using more detailed PG&E data, rested on the presumption that there may be a future need to communicate some of these results to the CPUC. Having retained much of the structure of the original CPUC model would assist in this communication.

However, once the Risk Team began building the model, it became clear that PG&E would require substantial modifications to the underlying model structure and user interface. The first adaptation required mapping existing program definitions into PG&E's unique segment categories. The Risk Model team created a matrix of all core, LGP and 3rd party programs that mapped each program's contribution into one of ten market segments.

The uncertainty model is actually a set of interlinked Excel spreadsheets. The primary original source of PG&E program and measure data is the Avoided Cost Workbook (v2d) (later versions of the model use MDSS data, the PG&E program-tracking database). Each program is modeled as a set of measures in a spreadsheet, with specific unit count, unit energy and NTGR data. General uncertainty data is contained in a separate, linked spreadsheet. This structure allows the user maximum flexibility in assigning global or local risk factors. The general goal is to identify the risk factors that have the largest impact on portfolio risk and replace global assumptions with segment, program or measure-specific data. The Risk Team researched the largest and riskiest programs and replaced global assumptions with specific data from recent research such as the California Measurement Advisory Council (CALMAC)[2] studies or work papers.

The long-term goal of this exercise is to continually update the model with the best data available, especially for those sectors in the portfolio where the model indicated large risks. For example, the first measure specific simulation was performed on CFLs in Mass Market Segment. The Risk Team identified this measure as having the largest potential impact on the portfolio, and expert judgment indicated that the existing risk factors, inherited from the CPUC activity, were not the most

[2] The California Measurement Advisory Council (CALMAC's) mission is to provide a forum for development, implementation, presentation, discussion, and review of market assessment and evaluation (MA&E) studies for energy programs within California. In the process of carrying out this mission, CALMAC coordinates and facilitates the dissemination and presentation of completed MA&E studies to stakeholders, policy makers and the general public. The lead representatives on the CALMAC are: 1) the Pacific Gas & Electric Company, 2) the California Energy Commission, 3) the California Public Utilities Commission Energy Division, 4) the California Public Utilities Commission Office of Ratepayer Advocates, 5) the Natural Resources Defense Council, 6) the Southern California Edison Company, and 7) Sempra (San Diego Gas & Electric Company, Southern California Gas Company).

recent available data. The Risk Team replaced the CPUC metrics with more recent unit count and NTGR estimates to create a new risk profile.

Creating the Dashboard

The model of PG&E's energy-efficiency portfolio blends thousands of assumptions. However, managers need to quickly and efficiently understand their portfolios and take action to improve their performance. Therefore, the Risk Team developed a user “risk dashboard" that summarized the portfolio using three key figures.

First, Figure 2 is an illustration of a possible distribution of results from the Monte Carlo simulation. Thus, based on the various ranges of assumptions within each program, this illustration suggests that PG&E's portfolio may save anywhere between 2.7 million and 3.4 million MWh. The median value was 3.1 million MWh. The probability that PG&E will reach or surpass their goal is represented by the lighter portion on the right of the histogram in Figure 2. In other words, 89 percent of all the simulated scenarios were generated a result where PG&E meet their savings goal. Conversely, there is an 11 percent chance (the red portion of the histogram) that PG&E will not achieve their goal based on the assumptions in this model.

Figure 2. Illustrative Data from Risk Dashboard Indicating the Probability of Achieving PG&E's Three-Year Goal for Their Overall Energy-Efficiency Portfolio.

Knowing the probability of achieving the goal is important, but does not alone suggest recommendations for improving such likelihood. Thus, Figure 3 illustrates the quantification of which segments contribute the greatest risk within the portfolio. Figure 3 suggests that over 95 percent of all the uncertainty in the portfolio is driven by the Mass Market Segment and Fabrication and Heavy Industrial Segment. CB calculated a contribution-to-variance for each of the nine segments that roll-up to form the entire portfolio. Figure 3 allows managers to see where to take action to reduce risk and thus improve the portfolio's performance.

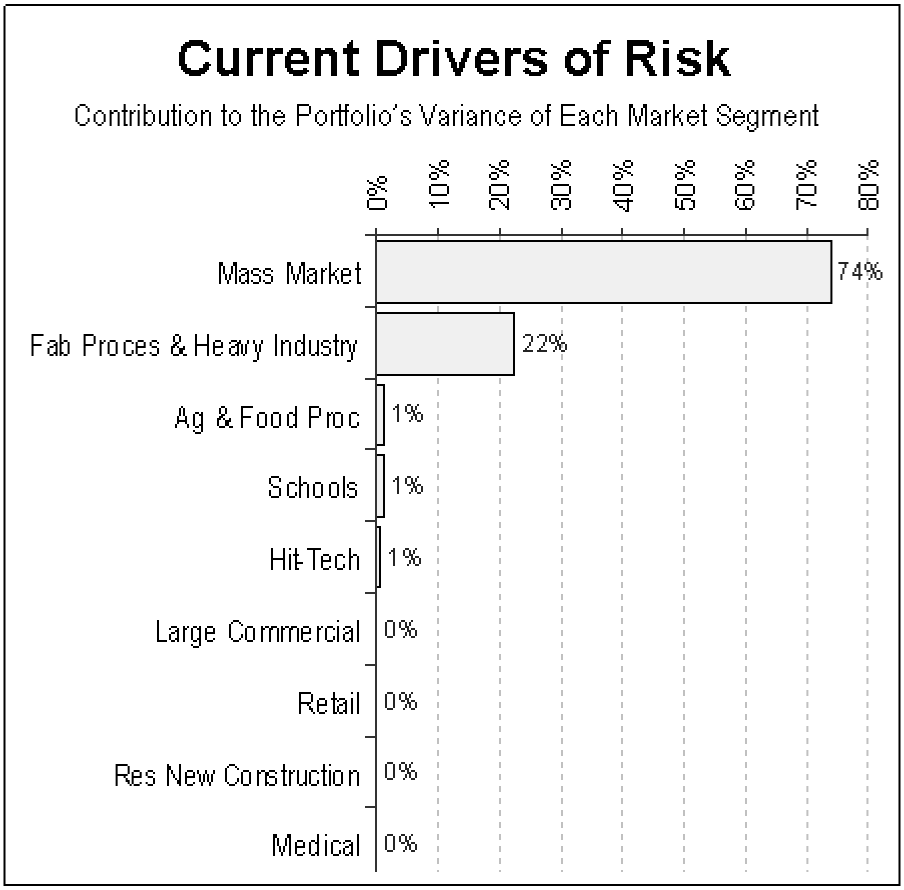

Figure 3. Illustrative Data from Risk Dashboard Indicating Which Segments Contribute to the Most Risk within the PG&E Overall Energy-Efficiency Portfolio.

The third and final figure on the Risk Dashboard is less intuitive and requires more explanation. The Risk Team continually changes the input assumptions in the portfolio model based on interviews with segment managers, data analysis, or secondary research. When such assumptions are changed, the simulation is re-run and generates a new distribution of results and sensitivities (similar to Figure 2 and Figure 3). A new simulation (or set of simulations) is referred to as an "Event". An Event is simply the results of the simulation based on a known set of input assumptions. It is interesting to see how the overall forecast changes with each new Event. Intuitively, one would assume that as the uncertainties around assumptions are reduced, the distribution of the results should also decrease. However, often, in the early phases of this process, the portfolio's uncertainty *increases* with the first few simulations because risks were either underestimated or not yet identified. Figure 4 tries to capture these dynamics by showing the forecasted energy savings and distribution of savings for five Events (e.g. five different simulations with different assumptions). The illustrative results in Figure 4 suggest that the new assumptions revised after Event #3 actually improved the chances of meeting the savings goal. This figure helps show the overall benefit of this risk management process. Simply stated, a systematic, data-driven, probabilistic approach to managing portfolio risk, will improve the chances for success.

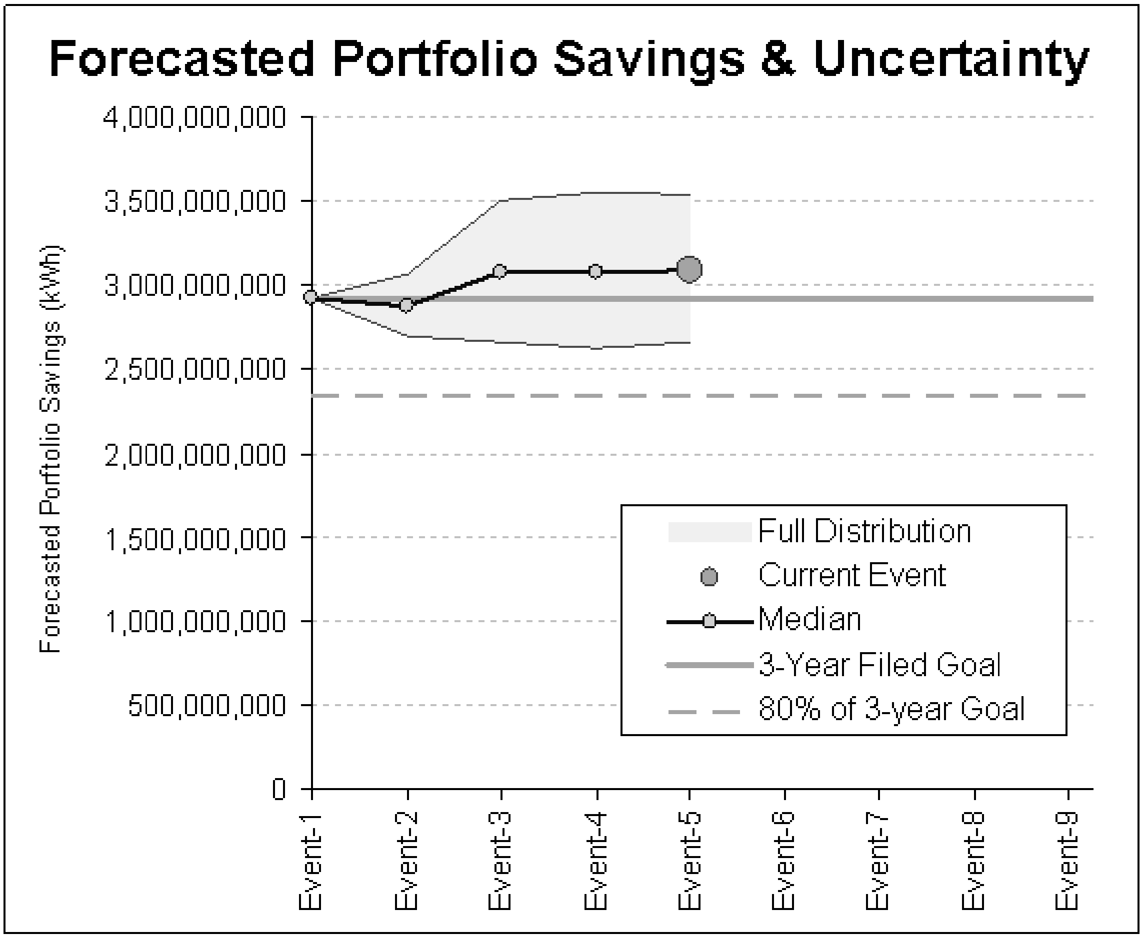

Figure 4. Illustrative Data from Risk Dashboard Indicating the Probabilistic Results from Five Different Portfolio Simulations Which Are Referred to as Events.

In addition to the three figures above, the dashboard also has a section summarizing the assumptions and changes included in each Event. One of the most important features of the dashboard is a dynamic interface allowing a user to "drill-down" into various programs, measure, or segments of the portfolio. For example, looking at Figure 3, the portfolio manager may see that most of the risk lies within the Mass Market Segment. Therefore, the portfolio manager can then look at the data for just this segment via a pull-down menu. When the user selects "Mass Market", all the figures change from showing data from the entire portfolio to showing only data from the Mass Market Segment. So, for example, Figure 2 would show the probability that the Mass Market Segment will reach its goal just as Figure 3 would show the primary programs (or measures) that drive the uncertainties. Once managers are able to identify and quantify risks, they can determine the most effective strategies to manage those risks.

Data Collection for Events

The initial risk model was based on PG&E's filing to the CPUC. These data, while sufficient as a general plan for meeting stated goals, were out of date with respect to current with program and segment implementation strategies. Therefore, the Risk Team interviewed each program manager to identify improvements to the model structure that would allow more effective risk modeling. The result is that each segment's risk profile now more closely represents the program manager's actual estimates. We give two brief examples. In the case of the Agriculture and Food Processing Segment, the program manager had already laid out goals based on previous year's experience with similar programs in relatively familiar markets. The likelihood of future program performance was based on the most recent three quarters of project results. For this segment, the Risk Team accessed PG&E's MDSS database and used the most recent data to adjust the original estimates in order to estimate future program performance. The second example comes from the Hi-Tech segment which contains some emerging technologies that, by definition, have little or no operating history. Thus, the likelihood of success for this segment is based on attaining a predicted percentage penetration of the market. This is an example of primarily expert judgment based risk.

It is the intent of the Risk Team to continuously update the model as new information on each segment becomes available. Each segment's relative contribution to total risk will be used to guide the application of future resources.

Using the Results

The purpose of performing a risk analysis is to identify, quantify and manage risks. The PG&E portfolio model indicates that there is substantial, manageable risk in the areas of unit count and NTGR. Some immediate activities are indicated. PG&E could commission internal evaluations to verify existing risks and potentially identify under-valued risk. Based on these results, PG&E could shift resources to better-performing and less risky program and measures and/or plan better programs for future cycles

Next Steps

The Risk Team has completed initial simulations of the PG&E energy efficiency portfolio. The current risk model reflects recent quarters program results, as well as updated expert judgments on expected market penetration in new segments. The model currently reports on kWh, kW and therms. Future activities include;

- Adding levelized costs to each segment
- Creating a new model using the latest version of the Avoided Cost Workbooks (3b5)
- Linking the model directly to MDSS and other PG&E databases.
- Using optimization features to perform portfolio optimization

Conclusion

Risk-based evaluation and assessment of utility energy efficiency programs is a growing and valuable service, both for the public sector evaluator and the private sector administrator. The techniques and tools described in this paper can enhance the reliability of energy efficiency implementation and increase the overall effectiveness of evaluation expenditures. By identifying the largest "pools" of uncertainty and risk, these tools allow all parties to identify, quantify and manage the current set of programs in order to improve future program performance.

References

Hall, Nick, Johna Roth, Carmen Best, Sharyn Barata, Pete Jacobs, Ken Keating, Ph.D., Steve Kromer, Lori Megdal, Ph.D., Jane Peters, Ph.D., Richard Ridge, Ph.D., Francis Trottier, and Ed Vine, Ph.D. 2006. *California Energy Efficiency Evaluation Protocols: Technical, Methodological, and Reporting Requirements for Evaluation Professionals.* Prepared for the California Public Utilities Commission.

Mathew, P., Kromer, S., Sezgen, O., Meyers, S. R., (2004), "Actuarial Pricing of Energy Efficiency Projects: Lessons Foul and Fair", Energy Policy.

Hall, Nick, Pete Jacobs, and Steve Kromer. 2006. *Improving the Reliability of Energy Efficiency Portfolio and Program Evaluation Via a Risk Assessment Approach*. Prepared for the Energy Division of the California Public Utilities Commission, 2006.

Violette, M., Freeman, R., (2004), "Demand Response Resources (DRR) Valuation and Market Analysis", Task XIII of International Energy Agency.

Schilmoeller, M., (2004), "Electric Power System Portfolio Risk Assessment", Proceedings of the 2004 Crystal Ball User Conference.