Risk Assessment & Management

This chapter presents the Council's approach to addressing uncertainty and managing risk. After reviewing the reasons for addressing uncertainty in the Council's Fifth Power Plan, it defines key terms and describes the analyses' principal sources of uncertainties. It describes how the studies evaluated the performance of resource plans under uncertainty, including their associated risk. The chapter introduces concepts that assist with the evaluation of cost-risk tradeoff, such as the "feasibility space" of plans and feasibility space's efficient frontier. It also describes the computer simulations the Council used to quantify plan performance and to construct the feasibility space. Most of the discussion of the computer model and the results of the simulations and analysis, however, appear in the next chapter. This chapter then briefly examines alternative measures of performance and risk and compares these to the selected approach.

In the last sections of this chapter, some background on and description of risk mitigation measures appears. Examples of risk mitigation measures are options to construct new power plants or implement voluntary load-curtailment programs. These sections emphasize the importance of planning flexibility in risk mitigation. Planning flexibility refers to the value of a planning option's ability to inexpensively and effectively respond to changing circumstances.

BACKGROUND AND ISSUES

The Western Electricity Crisis of 2000-2001 was a potent reminder that the electricity system is inherently risky. The crisis posed many important questions for the Fifth Power Plan:

- How much generation is enough and are there ways to assure its development?
- What is the value of demand response?
- What is the value of sustained investment in conservation?
- What is the value of resource diversity? How should uncertainty about fuel and wholesale power prices affect decisions about resource additions?
- How does transmission improve system reliability?
- What is the possible impact of global climate change on the power plan?

The evaluation of each of these issues depends on the decision maker's view of risk and uncertainty. For example, demand response is forthcoming if incentives exceed around $100-$150 per megawatt-hour. Even though demand response programs are relatively inexpensive to maintain, most forecasts of wholesale electric power prices rarely, if ever, exceed this value. The energy crisis of 2000-2001, however, demonstrated that unforeseen circumstances can send prices higher than this for extended periods. Consequently, key issues in this power plan require an analytical approach that addresses such rare but extreme events.

Risk assessment and management have always been important elements of the power plan. In prior plans, load uncertainty, fuel price uncertainty, and hydro generation variability figured prominently in the conclusions of the plan. Those plans incorporated gas and coal price excursions in forecasts and sensitivity analyses. They also considered capability to export and import various amounts of power to and from outside the region. Since the first power plan, the Council has analyzed the value of shorter lead times and rapid implementation of conservation and renewables. The Council has also valued "optioning" generating resources – carrying out pre-construction activities and then, if necessary, delaying construction until conditions favor going ahead.

In this plan, the Council further integrates risk assessment and management into its analysis and extends the assessment of risks to such issues as electricity market price uncertainty, aluminum price uncertainty, and emission control cost uncertainty. The analysis includes periods up to a few years when power and fuel prices, as well as other sources of uncertainty, deviate significantly from equilibrium levels. The study abandons the assumption of perfect foresight to better assess the value of risk mitigation.

DEFINITIONS

The following are terms used throughout the power plan.

- **Uncertainty** is a measurement of the quality of information about an event or outcome. Some future events are uncertain, but there is a significant amount of information about their likelihood. For example, the total annual flow at Bonneville Dam in 2010 is uncertain, but 61 years of historical records provide information about the distribution of outcomes. Other future events are less certain, like prices of natural gas and electricity. Theory and experience are informative to some degree, but expectations can be confounded. For others, there is very little information to go on. For example, there is almost no objective basis for determining the magnitude of any carbon tax in 2010. Future events therefore lie along a spectrum of varying degrees of uncertainty.

- **Futures** are uncontrollable events or circumstances. Futures are combinations of sources of uncertainty, usually specified over the entire 20-year study. For example, a future would include paths for loads, natural gas prices, water conditions, electricity market prices and so on over the 20-year planning period. Whether these sources of uncertainty produce risk or not depends on the adopted plan.

- **Plans** are future actions that are controllable. Plans include preparations to construct new power plants and the implementation of demand-side strategies or mechanisms. Power plant preparations include siting and licensing for specific construction start dates and quantities for each type of plant.. Different resources also provide differing amounts of planning and operating flexibility. These are inherent attributes of each plan.

- A **Scenario** is a plan considered under a specific future. The cost of a fixed plan under multiple futures provides the basis for cost distributions used in this study. Consequently, the costs of scenarios are fundamental to the study of uncertainty and risk.

- **Risk** is a measure of bad outcomes associated with a given plan. If the primary outcome of a study is the net present value cost over a study period, a bad outcome arises when a plan results in high development or use costs under a specific future. Risk is a measurement of the bad outcomes from the distribution of *all* outcomes associated with the plan under *all* the futures.

The Council has adopted a quantitative measure of risk. It captures both the likelihood and magnitude of bad outcomes. An unlikely outcome may still present significant risk if its effects are catastrophic.

- A **Risk Mitigation Action** is a plan, or some element of a plan, that reduces risk. A dispatchable power plant may protect a power system from high costs when electricity prices fall, because generation and fuel cost are curtailable. A demand-side strategy can protect a power system from fuel price risk and electricity market risk.

UNCERTAINTIES

What are some of the primary sources of uncertainty, and who bears the associated risks? What is the likelihood of particular futures and how do the various sources of risk conspire to produce particularly harsh futures?

The power plan addresses the following sources of risk:

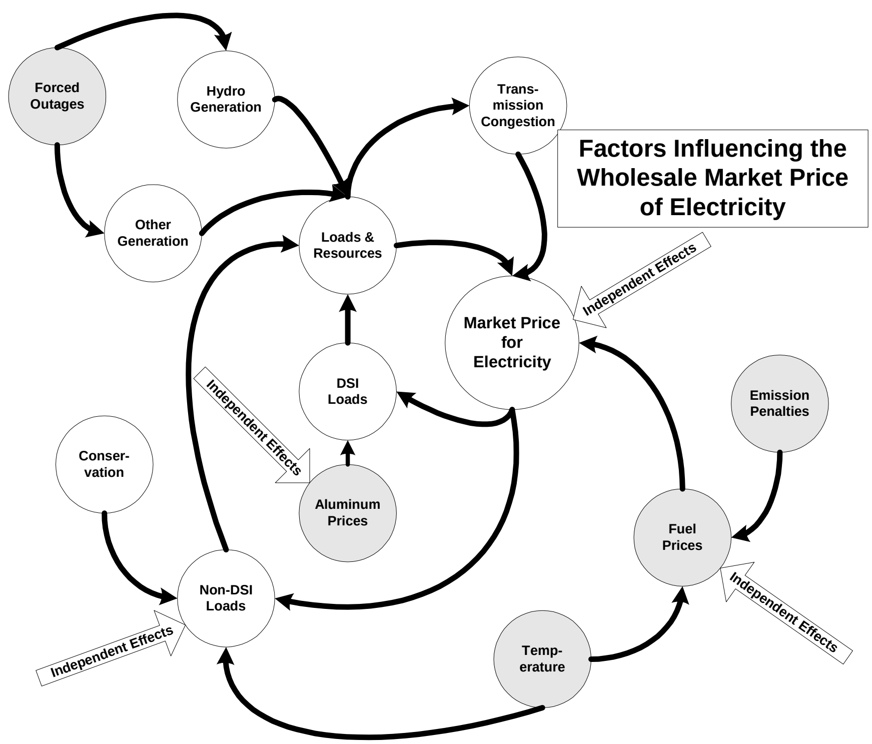

Figure 6-1: Factors Influencing the Wholesale Electricity Price

- **Wholesale Power Prices –** Many forecasters use long-term equilibrium models such as Aurora® to estimate future electric power prices. While useful to understanding price trends, these models ignore the disequilibrium between supply and demand that is commonplace for electricity. Disequilibrium results from less than perfect foresight about supply and demand, inactivity due to prior surplus, overreaction to prior shortages, and other factors. Periods of disequilibrium can last as long as it takes for new capacity to be constructed or released, or surplus capacity to be retired or "grown into." Resulting excursions from equilibrium prices can be large and are a significant source of uncertainty to electric power market participants. Because it is very difficult for an individual utility to exactly match loads and its own resources at all times, virtually all utilities participate in the wholesale market, directly or indirectly, as buyers and as sellers. This is particularly so when the region's primary source of generation, hydroelectricity, is highly variable from month to month and year to year.

To capture these effects, simulation models must vary electricity prices with hydropower availability, loads, and natural gas prices. The Council's portfolio model, described later in this chapter and in the next chapter, incorporates correlations among those factors. In addition, market prices must reflect changes in available generation relative to load. For a given load, additional generation tends to drive down electric power prices. In particular, if generation would initially exceed requirements, plus the region's ability to export, prices will be reduced until generation equals loads plus export capability. Similarly, if generation is inadequate to meet requirements, given the region's import capability, prices will increase until the situation is resolved, i.e., loads are reduced or the price induces sufficient generation.

Finally, electricity prices also exhibit substantial random variations due to conditions in other parts of the interconnected West and other factors that are not explicitly considered. These other factors include, for example, regulatory and legislative innovations and the introduction of new generation technologies. Figure 6-2 shows a sample of electricity price futures from the portfolio model.

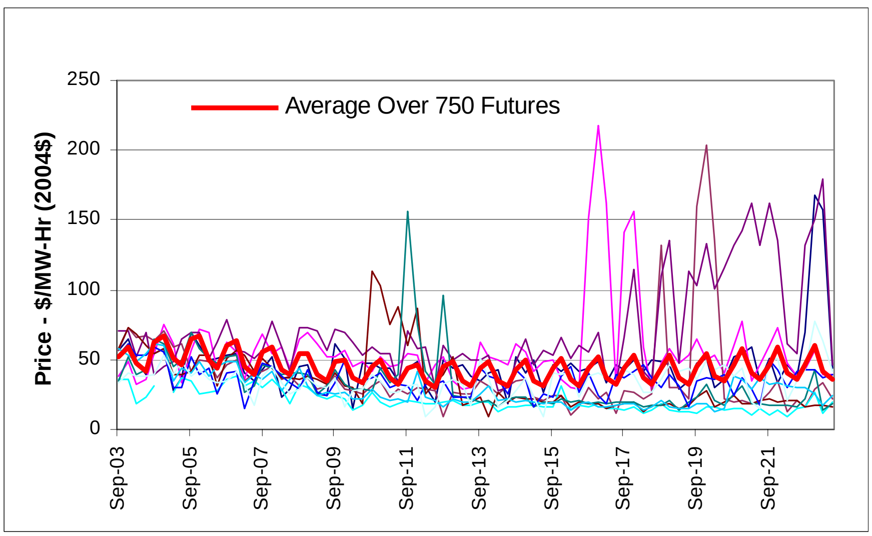

Figure 6-2: Electricity Price Futures

The Council contracted BHM3 Consultants to perform detailed statistical analysis on the relationships between hydro-generation, loads, temperature, natural gas prices, electric power prices, and transmission. The System Analysis Advisory Committee reviewed the results of these analyses. These analyses form the basis for the Council's representations of price paths, uncertainties, volatilities, and correlations. The results of these analyses are included in Appendix P.

- **Plant Availability** - Power plants are not perfectly reliable, and forced outages are an important source of uncertainty. The analysis includes simulation of forced outages based on typical forced outage rates for the generating technologies considered.

- **Load Uncertainty** - The Council's load forecast range for non-aluminum loads serves as a basis for the characterization of uncertain load trends. The expected load and the long-term load probability distribution are consistent with the forecast range. However, additional variations in load are added in the portfolio analysis to reflect seasonal and hourly patterns of load as well as excursions for weather variations and business cycles. Figure 6-3 displays a sample of load futures from the portfolio model simulations compared to the trend forecast range.

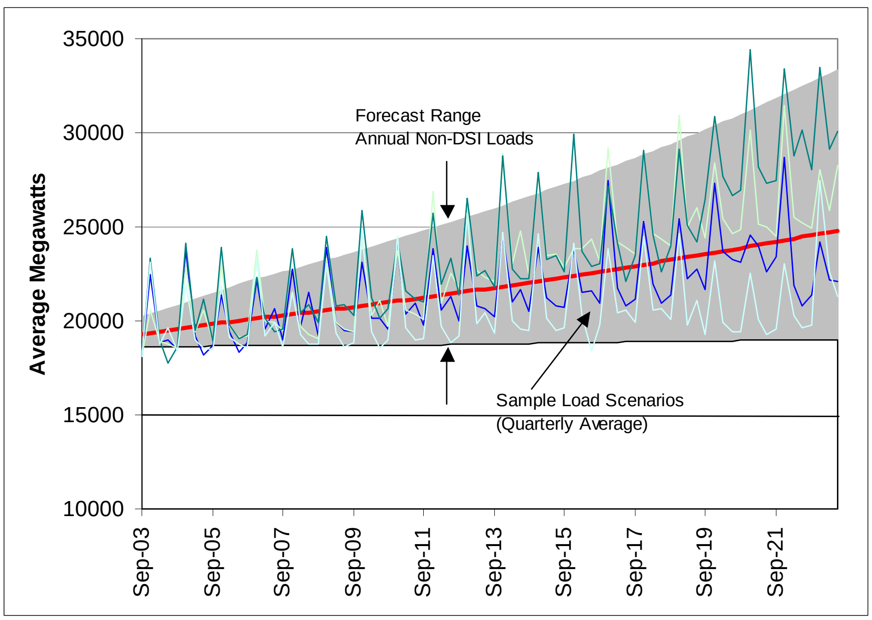

Figure 6-3: Load Futures

- **Aluminum Load Uncertainty** – Aluminum smelters in the Pacific Northwest have represented a substantial portion of regional loads in the past. This introduces a source of uncertainty directly related to the relative price of aluminum and the price of wholesale power. When electric power is costly relative to aluminum prices, smelters will shut down. The portfolio model captures the relationship among varying aluminum prices, electricity prices, and aluminum plant operation. In addition, the analysis considers the likelihood of permanent aluminum plant closure if a plant is out of operation for an extended period. Given the future electricity and aluminum price trends and variations and absent some policy intervention, the portfolio model results show a 80 percent likelihood of all aluminum plants closing during the forecast period.

- **Fuel Prices** - The basis for uncertain natural gas price trends is the Council's fuel price forecast range including estimates of uncertainty in the expected annual price. In addition to uncertainty in long-term trends in fuel prices, the modeling representation uses seasonal patterns and brief excursions from trends. These excursions may last from six months to four years and then recover back toward the trend path. The duration of the excursion and the duration of the price recovery are both functions of the size of the excursion. Figure 6-4 illustrates some natural gas price futures from the portfolio model simulations (2004$).

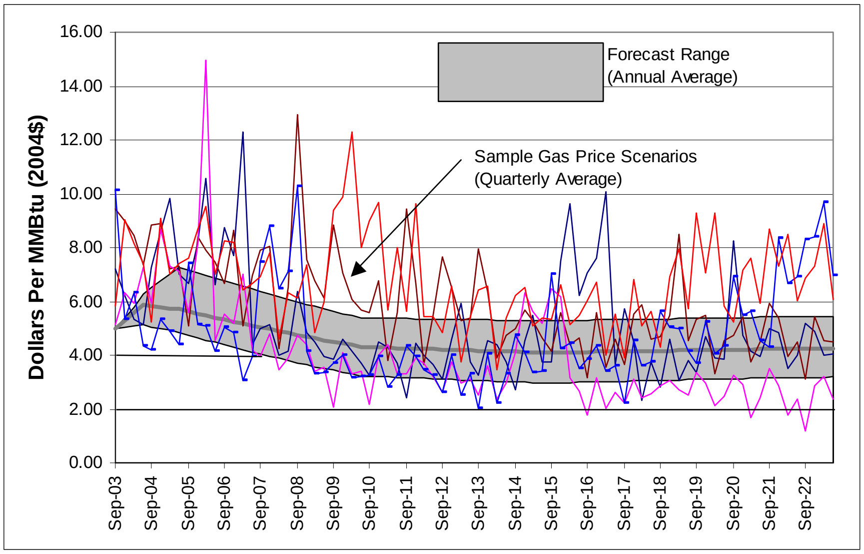

Figure 6-4: Gas Price Futures

- **Hydro generation** - A 50-year history of streamflows and generation provide the basis for hydro generation in the model. The hydro generation reflects constraints associated with the NOAA Fisheries 2000 biological opinion. The modeling assumes a decline of 300 average megawatts over the 20-year study period to capture relicensing losses, additional water withdrawals, the retirement of inefficient hydro generation units, and other factors that might lead to capability reduction. Hydro generation modeling did not reflect generation changes due to any climate change, because study results are too preliminary. Appendix N addresses work to understand any climate change impact on the hydroelectric system.

- **Climate Change -** A significant proportion of scientific opinion holds that the earth is warming due to atmospheric accumulation of greenhouse gasses. The increasing atmospheric concentration of these gasses appears to result largely from combustion of fossil fuels. Significant uncertainties remain, however, regarding the rate and ultimate magnitude of warming and its effects. The possible beneficial aspects to warming appear outweighed by adverse effects. A number of industrialized nations are taking action to limit the production of carbon dioxide and other greenhouse gasses. Within the United States, a number of states, including Washington and Oregon, have initiated efforts to control carbon dioxide production. It appears that the United States could eventually enact federal climate change policy involving carbon dioxide control. Further discussion of climate change policy appears in Appendix M

Because it is unlikely that reduction in carbon dioxide production can be achieved without cost, future climate control policy can be viewed as a cost risk to the power system of uncertain magnitude and timing. A cap and trade allowance system appears to be the most cost-effective approach to CO_2 control. The model, however, uses a fuel carbon content tax as a proxy for the cost of carbon dioxide control, whatever the means of implementation. The effect on existing power plant generation and the economic value of new generation would be representative of any type of effort to control CO_2 production using carbon-proportional constraints.

In the model, a carbon tax can arise in any election year. The probability of such a tax being enacted in at some time during the forecast period is sixty-seven percent. If enacted, the value for the carbon tax is selected from a uniform distribution between zero and $15 per ton if it is enacted between 2008 and 2016; and between zero and $30 per ton if enacted thereafter (2004$). Additional sensitivities are also considered.

- **Renewable energy production incentives** - Originally enacted as part of the 1992 Energy Policy Act to commercialize wind and certain biomass technologies, the production tax credit and its companion Renewable Energy Production Incentive have been repeatedly renewed and extended. These production tax credits (PTCs) have amounted to approximately $13 per megawatt hour on a levelized basis (2004$). The incentive expired in at the end of 2003 but, in September 2004, was extended to the end of 2005, retroactive to the beginning of 2004. In addition, in October, the scope of qualifying facilities was extended to include all forms of “open loop” biomass (bioresidues), geothermal, solar and certain other renewable resources that did not previously qualify. Though the amount and duration of the credit for wind remained as earlier, the credit for open loop biomass and other newly qualifying resources is half the amount available for wind and limited to the first five years of project operation. The longer-term fate of these incentives is uncertain. The original legislation contains a provision for phasing out the credit as the above-market cost of qualifying resources are reduced. Moreover, federal budget constraints may eventually force reduction or termination of the incentives. In the model, two events influence PTC value over the study period.

  The first event is termination due to cost-competitiveness. There is a small probability the PTC could disappear immediately, if congress decided renewable energy technology is sufficiently competitive and funds are needed elsewhere. The likelihood of termination peaks in the model when the fully allocated cost of wind approaches that of a combined cycle power plant around 2016. The probability falls to zero when the wind energy-cost forecast declines to 30 mills/kWh in 2034 (2004$). That is, there is never a modeling future where a PTC extends beyond 2034.

  The second event that modifies the PTC in the Council’s model is the advent of a carbon penalty. This event is related to the first, in that a carbon penalty would make renewables that do not emit carbon more competitive relative to those generation technologies that do. A CO_2 tax of less than about $15 per short ton of CO_2, however, would not completely offset the support of the PTC. For this reason, the value of the PTC subsequent to the introduction of a carbon penalty depends on the magnitude of the carbon penalty. If the carbon penalty is below half the initial value ($9.90 per megawatt hour in 2004$) of the PTC, the full value of the PTC remains[1]. If the carbon penalty exceeds the value of the PTC by one-half, the PTC disappears. Between 50 percent and 150 percent of the PTC value, the remaining PTC falls dollar for dollar with the increase in carbon penalty, so that the sum of the competitive assistance from PTC and the carbon penalty is constant at 150 percent of the initial PTC value over that range.

- **Green tags** - Power from renewable energy projects currently commands a market premium - a reflection of the perceived environmental, sustainability, and risk mitigation value of renewable energy resources. Driving the premium are above-market prices paid by utility customers for "green" power products, above-market prices paid for renewable energy components of utility supply portfolios and above-market prices for renewable acquisitions to meet requirements of renewable portfolio standards and system benefit charges. Tag value varies by resource and is reported to be between $3 to $4 per megawatt-hour for wind power, at present.

  In the model, green tag value can start the study period any where between $3 and $4 per megawatt-hour with equal likelihood (2004$). By the end of the study, the value can be anywhere between $1 and $8 per megawatt-hour (2004$). A straight line between the beginning and ending values determines the value for intervening periods. Consequently, green tag value averages 3.50 at the beginning of the study and averages $4.50 at the end of the study. Uncertainty in the value increases over time. This value is unaffected by events such as the emergence of a carbon penalty or the termination of the production tax credit.

- **Windpower shaping costs** - Windpower shaping costs are reported to range from $3-$8 per megawatt hour, lower than expected several years ago. The model uses deterministic shaping costs: $5.02 per megawatt hour for the first 2,500 megawatts of wind capacity and $10.76 per megawatt hour thereafter (2004$).

- **Other Emission Costs** - Power plant costs include the cost of the best available control technology required to meet current air emission requirements. The costs for coal-fired power plants also assume additional mercury control in anticipation of regulations currently under consideration by the Environmental Protection Agency.

- **Distribution Uncertainties and Modeling Errors** - An important source of concern to decision makers is the validity of a computer model's representation, the accuracy and completeness of input data, and the potential that a user may simply make a mistake in applying the model.

  One of the mechanisms for dealing with this sort of risk is a careful evaluation of whatever plan is produced by the computer model. Regardless of the nature of the uncertainties and the probabilities associated with futures, the resulting plan must make sense to the decision maker, and the means of risk mitigation must be clear and

[1] The conversion of carbon penalty ($/US short ton of CO_2) to $/MWh is achieved with a conversion ratio 1.28 #CO_2/kWh. This conversion ratio corresponds to a gas turbine with a heat rate of 9000 BTU/kWh.

compelling. The Council uses models to screen plans, not as a substitute for experience and judgment.

Sensitivity analyses around the distributions for the uncertainties provide another check on the modeling work. The resource plan produced by Council staff incorporates distributions of forecasts prepared or reviewed by experts in an open forum. Although "uncertainty about uncertainty" does not make sense for a single decision maker, there nevertheless will be a diversity of opinion among decision makers about the uncertainty of specific forecasts.

The Council's model does not explicitly treat other sources of uncertainty, such as changes in technology and policy, fish and wildlife programs, and the transmission system. Appendix P describes these sources of uncertainty and any treatment in this analysis.

OBJECTIVES AND MEASURES OF RISK

How are the costs and benefits of plans determined? How is risk measured? Whose assessment of risk should be used?

The Council's approach to resource planning is called "risk-constrained least-cost planning." Given any level of risk tolerance, there should be a least-cost way to achieve that level of risk protection. The purpose of the Council's analysis is to define those plans that do just that.

Given a particular future, the primary measure of a plan is its net-present value total system costs. These costs include all variable costs, such as those for fuel, variable operation and maintenance (O&M), certain short-term purchases, and fixed costs associated with future capital investment and O&M. The present value calculation discounts future costs to constant 2004 dollars using a real discount rate of four percent.[2] This treats current and future costs on a comparable basis. Total net-present value costs are demonstrably a better measure of economic value than internal rate of return, retail power rates, or benefit-cost ratio.

If the future were certain, net present value system cost would be the only measure of a plan's performance. Because the future is uncertain, however, it is necessary to evaluate plans over a large number of possible futures. Complete characterization of the plan under uncertainty would require capturing the distribution of outcomes over all futures, as illustrated in Figure 6-5 below. Each box in Figure 6-5 represents the net present value cost for a scenario sorted into "bins." Each bin is a narrow range of net present value total system costs. A scenario is a plan under one particular future.

Because a simulation typically uses 750 futures, the resulting distributions can be complicated. Representative statistics make manageable the task of capturing the nature of a complex distribution. The *expected* net present value total system cost captures the central tendency of the distribution. The expected net present value is the average of net present value total system costs, where the average is frequency weighted over futures. This plan will often use the shorthand expression, "average cost of the plan." The average cost is identified in Figure 6-5.

[2] See Appendix L.

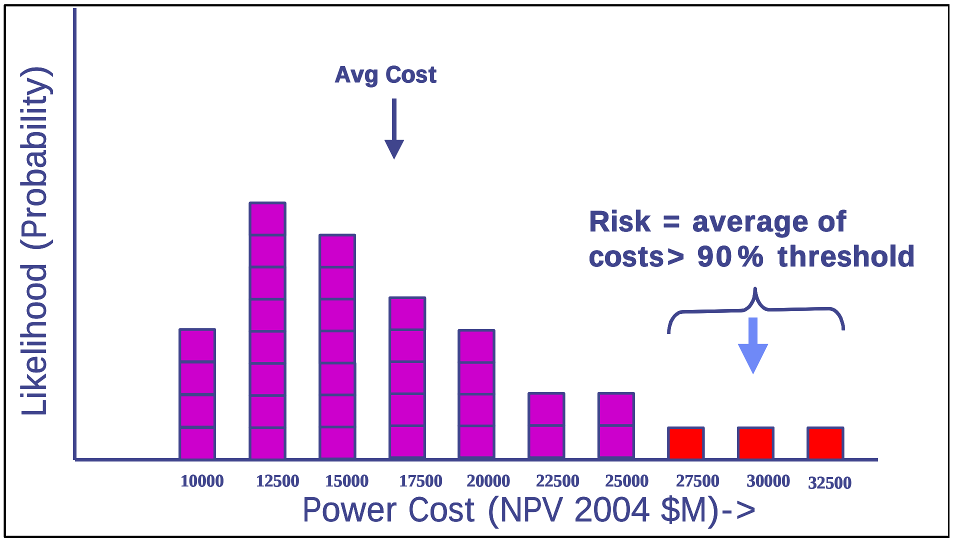

Figure 6-5: Representative Plan Cost Distribution, Average Cost and Risk (TailVaR$_{90}$)

Expected net present value cost, however, does not give a picture of the risk associated with the plan. There are a number of possible risk measures that could be used. A summary measure of risk called "TailVaR$_{90}$" was chosen. This choice of risk measure and its comparison with other risk measures appears in Appendix P. Very briefly, TailVaR$_{90}$ is the average value for the worst 10 percent of outcomes. It belongs to the class of "coherent" risk measures that possess mathematical properties superior to alternative risk measures. Since 1998, when papers on coherent measures first appeared, the actuarial and insurance industries have moved to adopt these, abandoning non-coherent measures such as standard deviation and Value at Risk (VaR).

ASSESSING PLAN PERFORMANCE AND IDENTIFYING OPTIMAL PLANS

The primary tool for identifying risk-constrained, least-cost plans is an analytical system with three components. The first component is Olivia, which creates an Excel® workbook portfolio model. With minor refinements, an early version of this workbook model has served as the regional model. Olivia is part of an effort that extends beyond the fifth regional power plan. The vision is to provide data and tools like Olivia to others -- such as utilities and public utility commissions -- to help them perform their own risk analysis using concepts and techniques developed by the Council.

The second component is the Excel workbook model itself, "the portfolio model." This workbook model is the calculation engine. It estimates costs of generation, of purchases and sales of wholesale power, and of capacity expansion over the 20-year study period. An Excel add-in runs a Monte Carlo simulation of the scenarios, with each game corresponding to a future.[3] This simulation gives rise to the cost distribution illustrated in Figure 6-5 for each plan.

Figure 6-6 illustrates the kind of calculation that the portfolio model makes in a specific scenario. It shows energy use resulting from a plan over a two-year time period for the fixed future. The future defines hydro generation, loads, gas prices, and so forth in each hour. Given these

[3] Decisioneering's Crystal Ball®. Olivia produces a workbook that is compatible with Crystal Ball.

circumstances, existing and future resources in the plan generate power, largely in response to wholesale electricity prices. Because generation rarely exactly matches load, power is purchased from or sold into the wholesale market. The model sums costs and revenues in each hour, adds any future fixed costs for existing and new generation or capital costs for new generation and conservation, and discounts the dollars to the beginning of the study. Of course, the portfolio model uses 20 years of costs, not two years, but the process is identical.

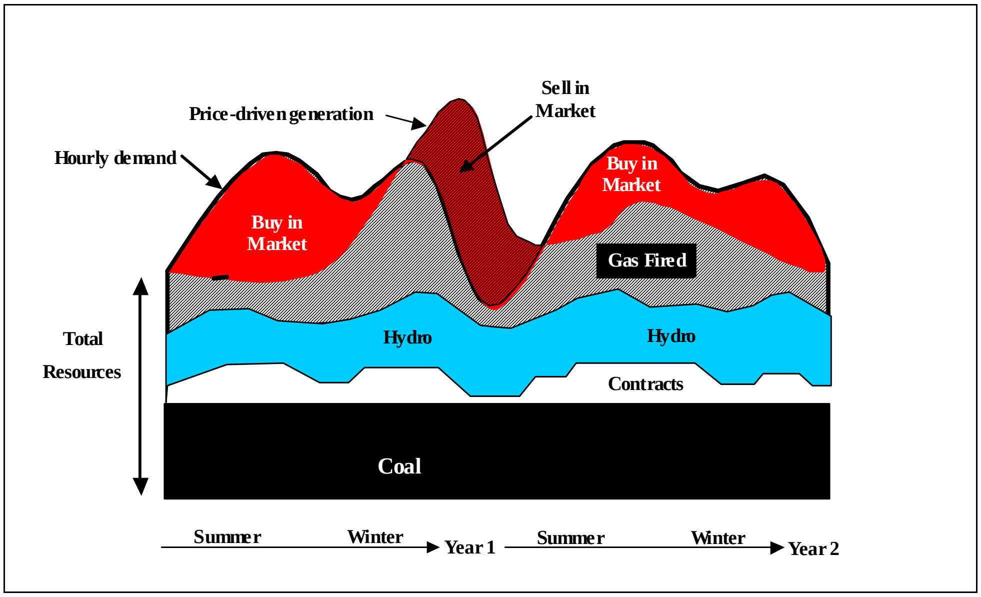

Figure 6-6: Portfolio Model Calculation

Typically, simulating 750 futures for each of 1400 plans would require that around a million scenarios be examined. If hourly calculations were performed for each of these 20-year scenarios, computation time would be prohibitive.[4] For this reason, algorithms were developed to estimate plant capacity factors, generation, and costs for periods of one to several months. Using these techniques, the 20-year study period is represented by 80 hydro-year quarters on peak and another 80 off peak. Since the model does not break the Northwest into sub-regions, cross-Cascade and other intra-regional transmission constraints are not modeled, but imports and exports are constrained to 6,000 megawatt-quarters, before any contracts.[5] Transmission constraints within the region are considered outside the model. Existing regional thermal resources are aggregated down to about 30 plants with similar characteristics. Hydro generation is based on draws from a 50-year streamflow record and system constraints determined by the 2000 Biological Opinion (BiOp). Operation of the region's seven remaining smelters is determined by the relative price of aluminum and wholesale electricity.

The third component of the analytical system helped find the least-cost plan for a given level of risk. This component is actually another Excel add-in.[6] This add-in uses a variety of techniques to find the least-cost plan for a given level of risk as efficiently as possible. The process of selecting a risk constrained least-cost plan is illustrated with the following diagram:

---

[4] One estimate using AURORA® run times put the study at a little over 85 years.
[5] Contracts may be fully counter-scheduled.
[6] Decisioneering's OptQuest®.

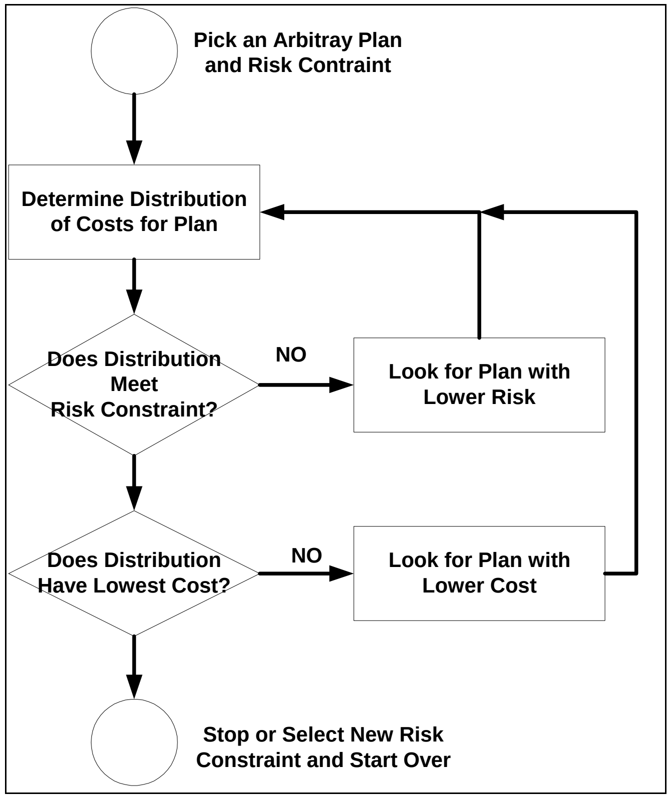

Figure 6-7: Finding the Risk-Constrained Least-Cost Plan

The program first seeks a plan that satisfies a risk constraint level. Once it finds such a plan, the program then switches mode and seeks plans with the same risk but lower cost. The process ends when a least-cost plan for each level of risk is found.

The necessity of using the approach to find the least-cost plans becomes evident when one attempts to estimate the number of potential plans that may exist. Assume that cumulative capacity expansion for four or five resource candidates were constrained to half a dozen levels at each of eight points in time.[7] Even with this modest choice, the number of potential plans is billions of billions.

If the outcome for each plan is plotted as a point with coordinates corresponding to the expected cost and risk of the plan, one obtains the new distribution illustrated in Figure 6-8. Each point on the figure represents the average cost and TailVar_{90} value for a particular plan over all futures. The least-cost outcome for each level of risk falls on the left edge of the distribution in the figure. The combination of all such least-cost outcomes is called the "efficient frontier." Each outcome on the efficient frontier is preferable to the outcomes to the right of it, since it has the same risk as those outcomes, but lowest cost. Choosing from among the outcomes on the efficient frontier, however, requires accepting more risk in exchange for lower cost, or vice

[7] The exact choice of the points in time differ from study to study, typically more of these points are concentrated around the most interesting periods, like the near-term action plan horizon or when a lot of resource expansion is taking place.

versa. The "best" outcome on the efficient frontier depends on the risk that can be accepted. This topic is described in greater detail in Appendix P.

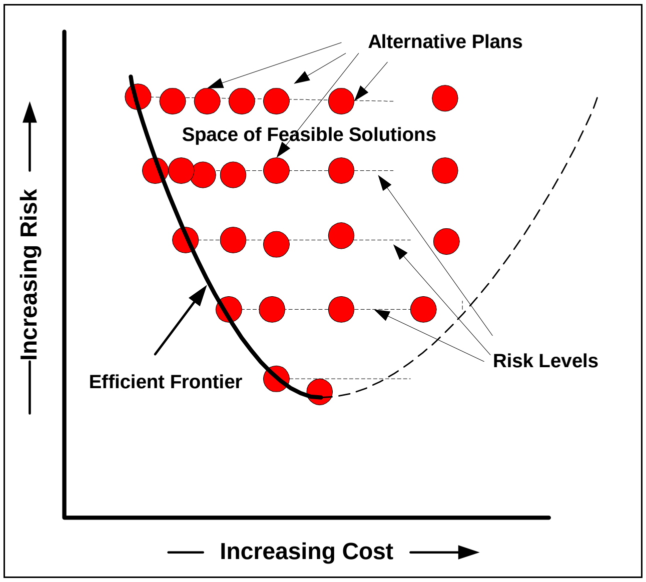

Figure 6-8: Feasibility Space

Application of the Efficient Frontier

One of the difficulties that surfaces when valuing resources is how to normalize the value for differences in risk mitigation. For example, one decision maker may adopt a least-cost plan that provides no protection from fuel price risk because he or his constituents have a high tolerance for risk. He may not be willing to pay a premium for a resource like wind generation that protects his constituents from excursions in rates due to fuel price volatility. Another decision maker may prefer a low-risk plan, which requires paying some premium -- relative to today's view of cost effectiveness -- to acquire those resources that provide this protection. Consequently, value typically follows risk aversion.

The efficient frontier provides a means to quantify the value of resources that provide some kind of risk mitigation and to explicitly describe that value as a function of risk aversion. Removing the resource as an option for capacity expansion leads to one of two outcomes: either the efficient frontier shifts to the right[8] along part of its extent, indicating costs have increased, or it does not. If the efficient frontier does not shift, it means other resources are capable of substituting for the resource in question. The resource therefore has no value at that level of risk. If the frontier shifts, the difference in cost at each risk level is the risk value of the resource. Figure 6-9 illustrates this effect for a typical resource.

[8] Removing the resource creates a set of resource options for the cost-minimizing logic that is strictly smaller than the original set. This implies that costs may not decrease.

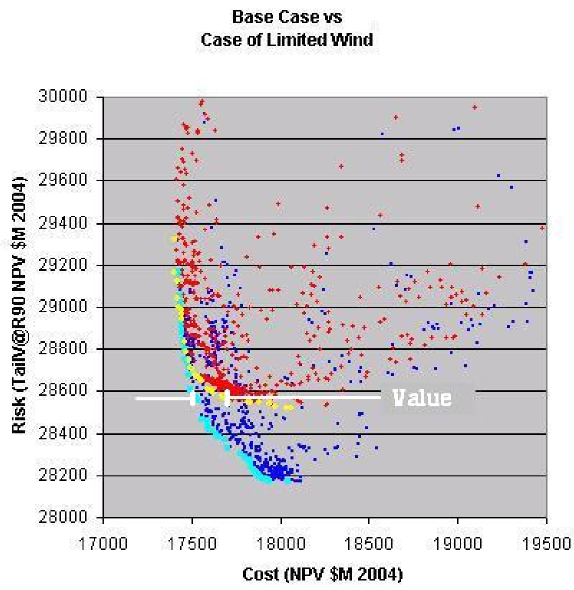

Figure 6-9: Value of Risk Mitigation Resource

What does this value represent? The efficient frontier using the candidate resource already captures the costs for planning and construction of the resource. The difference represents the value associated with assuring the resource is available at the cost and terms the model assumes.

The Risk/Cost Trade-off

The Council and the System Analysis Advisory Committee have extensively discussed the issue of risk/cost trade-off. A single choice cannot represent every decision maker in the region. Arguably, it may be meaningless to attempt to arrive at such a risk/cost trade-off for the region. While it may not be possible to settle on a level of risk tolerance that represents all parties in the region, consideration of risk issues and the efficient frontier can provide insights for the Council and others in the region.

First, it may not be necessary to pick a risk/cost trade-off. The efficient frontier alone can yield significant insights. Attributes that are common or absent from among all the plans on the efficient frontier can help a decision maker to identify robust resource strategies and flag potential strategic blunders. As the next chapter explains, the common attributes turned out to be key to elements of the Council's Action Plan.

Second, many plans along the efficient frontier may differ only by commitments that do not need to be made today. If the earliest resource commitments from among all the plans occur at some point in the future, decision makers can and should wait until to make them. At that future time, the decision makers will have more information and a better choice may be more apparent. A useful product of the exercise, nevertheless, is observing the earliest that any commitment would be necessary. This observation can inform the timing of future planning efforts.

Third, partitioning plans along the frontier into classes of strategy can make planning more manageable. Typically, plans along the efficient frontier do not follow a simple, "more resources mean less risk" pattern. The analyst will observe regimes where different technologies and strategies prevail, or different kinds of risk dominate. Using representative plans from each regime can help simplify subsequent analysis.

Fourth, the region's risk situation is likely to be representative of that which some parties in the Pacific Northwest power industry face. The conclusions and methods presented here may help others with assessing and communicating their risks and risk mitigation activities. Several general principles, described in the next chapter, may have value to individual industry participants.

Moreover, the analysis presented in the power plan identifies a value for risk mitigation resources and programs to the region. Focusing exclusively on the least-cost plans without consideration of risk could expose the region to significant risk. As described in the previous section, this analysis estimates the value of various resources and actions, including risk mitigation value.

At the beginning of this chapter, the question was raised, "How much generation is enough, and are there ways to assure its development?" Regulatory and legislative policies may be necessary. However, power systems and financial systems are complex, and regulatory policies can create ill-fitting, inequitable, and inefficient solutions in any particular situation. Individual participants, by insuring themselves against the risks that they face can help to secure a more reliable system for the region. The tools and methods developed for this analysis are available to decision makers in the region. Standard reports of risk analysis results may help regulators, load-serving entities, and their constituents communicate better.

The next chapter will take a more in-depth look at the issue of risk/cost trade-off for the specific efficient frontier the planning process delivered and plans along that frontier. TailVaR_{90} has served well as a means of screening plans for risk, but that chapter presents values for other kinds of risk. This issue of alternative measures for risk raises the question of how TailVaR_{90} compares to alternative risk measures.

ALTERNATIVE MEASURES OF COST AND RISK

Many alternatives exist for measuring the central tendency and the risks associated with a distribution of costs. Would the efficient frontier look different using alternative measures of risk or cost?

The results of portfolio model studies include a host of alternative measures. For each feasibility space study, data for each plan include both the median and the mean cost of the distribution. The model also tracks a host of risk measures. Risk data include:

1. TailVaR_{90}
2. Standard deviation
3. VaR_{90}
4. CVaR_{20000}
5. 90th Decile
6. Mean (over futures) of maximum (over 20 years) of annual cost increases
7. Mean (over futures) of standard deviation (over 20 years) of annual costs

Subsequent studies examined alternative sources of risk, such as relative exposure to bad market conditions and variation in average power cost, including embedded costs.

Mean costs and $TailVaR_{90}$ do a reasonable job of screening plans. For modeling the regional portfolio, there is a strong consistency between the chosen measures and the alternatives in most cases. This correspondence is neither accidental nor universal among load-serving entities in the region. This section describes the relationship for a couple of alternative measures and provides a reason for the correspondence. A complete treatment of the alternative measures appears in Appendix P.

An Alternative to the Mean

Some would argue that the median is a better measure of central tendency than the mean for risk analysis. The median future is a future above and below which lie an equal number of better and worse futures. In contrast, a weighing scheme defines the mean: the mean is the average of outcomes, weighed by their probabilities.   What future will the region face? For that matter, what determines the outcome of rolling dice? It is a matter of the likelihood of landing on each face, not the value of the faces. The mean cost, in fact, may not correspond to any particular future, just as there is no face on a die with the value 3.5, the average outcome. For an odd number of futures, however, there is always a median value future[9].

On the other hand, the mean is a statistic with which most decision makers seem to have greater comfort. Some decision makers may feel that they want extreme outcomes to influence their measure of the central tendency.

Fortunately, it does not seem to make much difference to an analysis of regional risk. Distributions for outcomes of plans exhibit a strong relationship between the two measures. Figure 6-10 shows that the mean and median values track very closely.

The mean value is consistently above the median, suggesting that distributions of cost are skewed. The distributions have long tails extending in the high-cost direction, pulling up the mean. As costs go down, the skewing becomes more pronounced. This has implications to the discussion of risk measures.

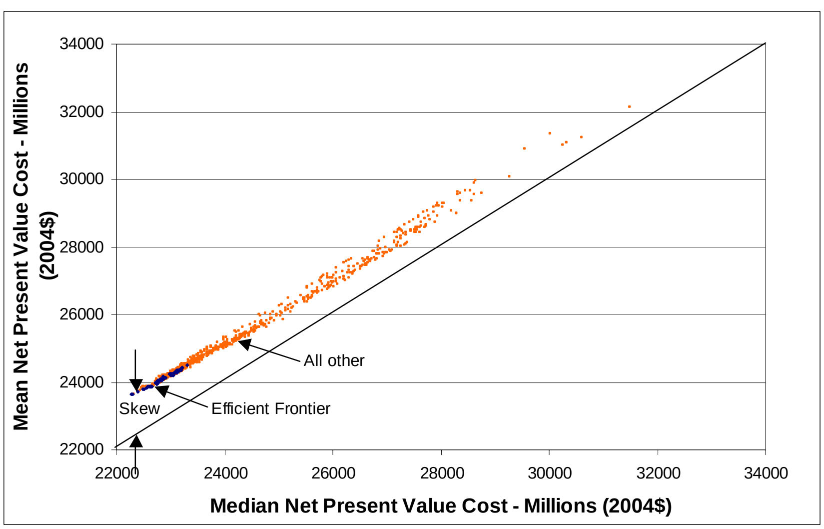

Figure 6-10: Mean vs Median

[9] The median of an even number of observations is the arithmetic average of the two middle observations.

Distributions of Cost for Regional Risk Study

Distributions of cost for typical load-serving entities or generators in the region differ significantly from that of the region as a whole, because individual participants are usually price takers. That is, their individual loads and the operation of their resources typically will not move prices in the region. If they have surplus resources, in particular, their potential for making money is large. It depends only on how the market price for electricity goes.

One potential cost distribution situation for price takers with surplus resources appears in Figure 6-11. Market prices for electricity may go down significantly, but prices are ultimately limited in how far down they can go. Costs are therefore limited to the fixed costs of resources surplus to requirements. The costs have a tail extending to the left in Figure 6-11 because high prices produce revenues that offset these fixed costs. Net costs can even become negative if prices, which are unbounded above, are high enough.

This situation does not arise for the region. This is because the aggregate regional resource situation affects market prices. Surplus resources for the region depress price. To see this, start with a resource-deficit situation. The cost distribution is skewed in the opposite direction (Figure 6-12), with a tail to the right. The reason for this shape is similar to that for the resource-surplus price taker, except that now there is a floor on fixed cost and more exposure to high prices of market energy to make up meet requirements. (Price takers who are *deficit* resources will potentially have this distribution of costs, as well.) The more resources added, the shorter the tail to the right becomes because the net requirement met by the market diminishes. When the resources are roughly in balance with the loads, the distribution becomes nearly symmetric. However, surplus regional resources depress price, so the distribution does not “flip around.” Export constraints limit the sales from any surplus resources. Once the export constraints become binding, prices fall and so do profits from sales.

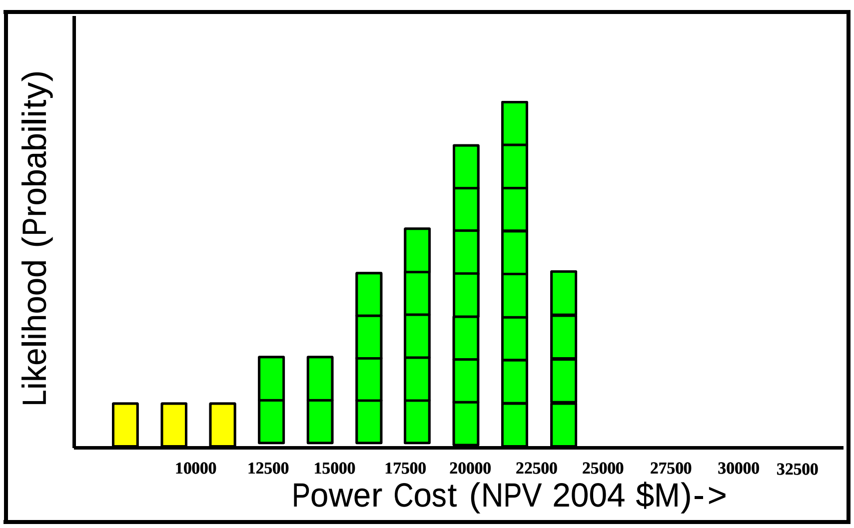

Figure 6-11: Cost Distribution for Resource-Surplus Price Taker

Because distributions like that in Figure 6-11 never arise in the regional study, the median cost is always lower than the mean cost. Moreover, what typically occurs is that the least-cost, highest-risk plan consists of relying on the market to meet requirements. In this case, of course, the distribution for regional costs becomes highly skewed. This explains why skewing becomes more pronounced in Figure 6-10 at the lowest average cost.

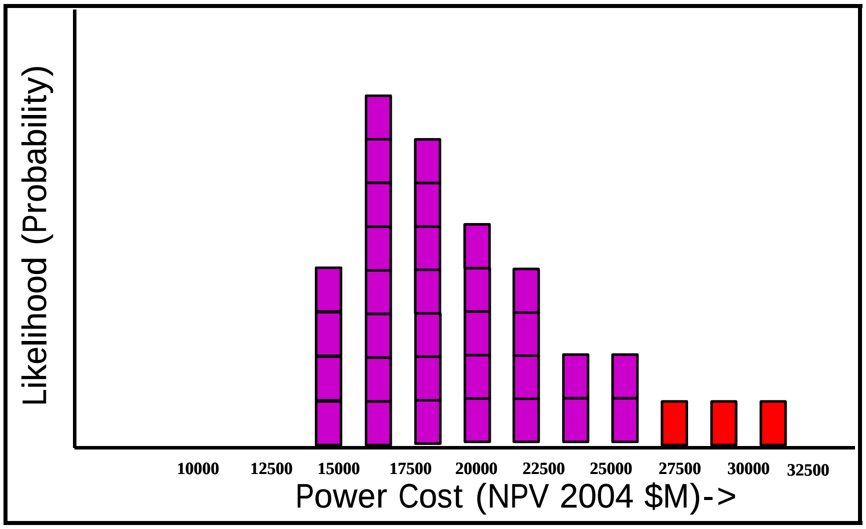

Figure 6-12: Cost Distribution for Region

Alternatives to TailVaR_{90}

The Council's portfolio analysis suggests that TailVaR_{90} is a reasonable risk measure for the region. This section will explore alternative risk measures and explain why TailVaR_{90} provides good guidance in evaluating regional plan risk.

To understand why TailVaR_{90} is robust, consider the example of standard deviation. Figure 6-13 restates the feasibility space associated with the base case,[10] using standard deviation as the risk measure. In Figure 6-13, the white points are the plans in the efficient frontier or near-efficient frontier using TailVaR_{90}. Clearly, these are also efficient using standard deviation. The black diamonds are the plans that are efficient using standard deviation, but not efficient using TailVaR_{90}. These require explanation. The smaller, interior points correspond to plans that are not efficient using either risk measure.

To understand what is taking place here, recall the description of the distribution of costs for the region in the previous section. Cost distributions for the region are skewed in only one direction. If a plan is on the efficient frontier using TailVaR_{90}, the cost is minimal so the distribution is typically as narrow as possible given the level of risk. That is, a plan with a narrower distribution would have higher cost and would not be least cost. Standard deviation is therefore as small as possible. The converse is not true, in general. A plan could represent a lot of resource addition, which would suppress prices and create a very predictable, but expensive outcome, that is, a narrow distribution of costs. This is evident in the plans represented by the black diamonds. It is unlikely that a decision maker using standard deviation would choose these plans, however. In these plans, the cost for *each future* is higher than the cost any of the plans represented by the white points *under the same future*. This is not evident from the graph, but it is one of the important properties of coherent measures of risk, like TailVaR_{90}. Supposedly, a

[10] See description and discussion in the next chapter.

rational decision maker using standard deviation could identify these plans by their high cost and would exclude them from consideration.

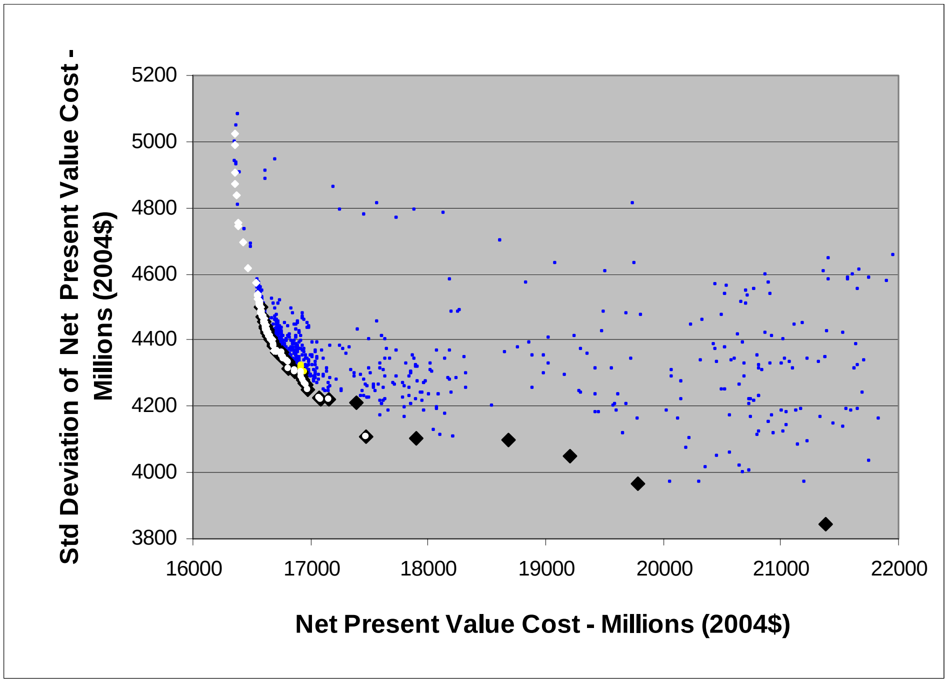

Figure 6-13: Using Standard Deviation to Create the Feasibility Space and Efficient Frontier

The analysis of correspondence of $TailVaR_{90}$ to other risk measures appears in Appendix P. Other coherent measures, like $CVaR_{20000}$, generally show a direct correspondence to $TailVaR_{90}$, whereas non-coherent measures, like Value at Risk (VaR), have a relationship to $TailVaR_{90}$ that resembles that of standard deviation. Along the efficient frontier, non-coherent measures will agree with $TailVaR_{90}$; away from the frontier, the agreement will be weak, but higher average cost would tend to remove the plans from consideration. The Council concluded that $TailVaR_{90}$ provides a robust measure of risk for screening regional plans.

Although cost distributions for the region skew in one direction, this is not true of load-serving entities that are price takers in the region. For these entities, coherent measures of risk are a better choice than non-coherent measures. The fortunate correspondence that exists between coherent and non-coherent measures of the regional cost distribution does not exist in their situation. For example, given a choice between a plan that produces a surplus and a plan that produces a deficit, the situation may arise where average cost and standard deviation are identical, as illustrated in Figure 6-14. The risk situation, however, is quite different. The deficit plan has a tail that extends in the high-cost direction, a riskier situation.

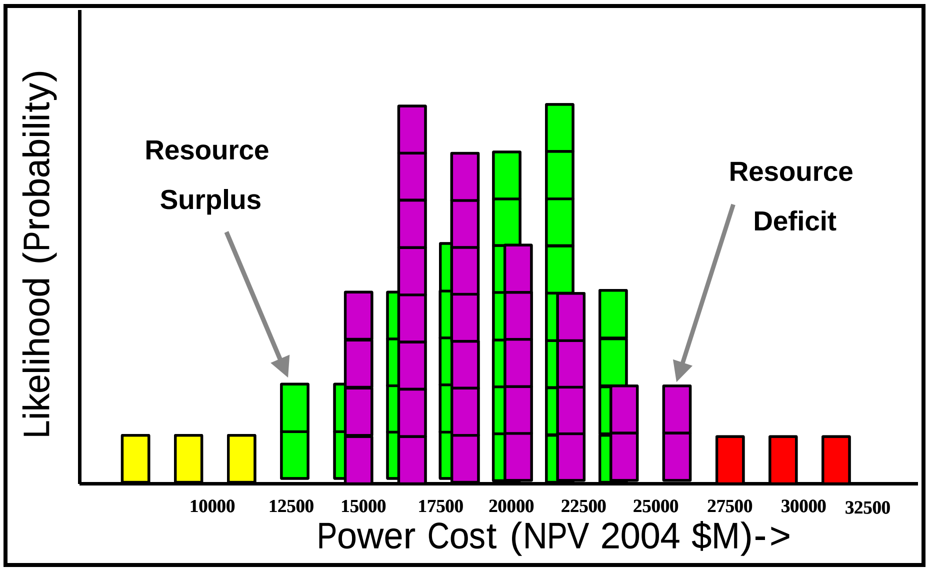

Figure 6-14: Price Taker's Cost Distributions for a Resource-Surplus Plan and a Resource-Deficit Plan

This concludes the discussion of uncertainties and sources of risk. This chapter now turns to the topic of risk management and mitigation.

RISK MITIGATION ACTIONS

How risk reduced, and what is the right amount of risk reduction? What level of risk mitigation does the region require? How do various resources contribute to risk mitigation? Who pays for risk mitigation?

The value of risk management resources is the contribution they make when foresight is not perfect. Their value derives from their ability to respond under abnormal circumstances of price, loads, resource availability, and so forth. Moreover, their value is directly related to the probability of these events.

Risk mitigation can be thought of as resulting from two types of actions:

- Hedging - A commitment to a plan that symmetrically reduces uncertainty; and
- Flexibility or optionality -- the right, but not the obligation, to take a particular action

Hedging

"Hedging your bets" is a common phrase. It means an action that will offset the effects of another action. In the specific context of power planning, a utility may want to add wind generation to its resource portfolio if the portfolio contains a lot of combustion turbine generation and the utility is concerned about risks of natural gas price increases. If natural gas prices go up, the utility's costs do not go up as much on average as they would if it had not invested in wind generation. However, hedges are not free. If natural gas prices decrease, some of the reduction in natural gas costs is offset by the utility's commitment to the fixed costs of a wind power plant. By this definition, the effect of hedging is always symmetrical, mitigating the worst outcome, but moderating the best outcome as well.

Flexibility or Optionality

In contrast to hedging, there are risk measures that provide “optionality” or “flexibility” and are asymmetrical in their effect. One of the most familiar examples is insurance. Home and car insurance protects the owner from futures or situations that diminish or wipe out the value of the investment. The insurance premium offsets a part of the value of the investment in all futures, but the insurance shields from loss in bad futures.

There are a host of examples of options in the power industry. The Council’s plan deals primarily with physical processes or decision-making flexibility. Because of the strong association with financial options and the confusion that association may create in the reader’s mind, the term “flexibility” will be used when referring to physical process or decision-making flexibility.

Examples of short-term flexibility are plentiful. A combustion turbine, for example, represents the flexibility to exchange natural gas for electric power. When electricity is expensive relative to natural gas, the turbine’s owner tends to sell the electricity generated from the gas. If natural gas is expensive relative to electricity, the owner tends to refrain from generating and may either resell the valuable gas or hold it in storage. Demand response represents another form of this flexibility. When electricity is expensive relative to a commodity that a utility customer is producing, the load serving entity and its customer may agree to sell the more expensive electricity and compensate the customer with more money than the customer would have made producing the commodity. These are examples of short-term flexibility.

Examples of long-term flexibility include a decision maker’s ability to cost-effectively cancel or defer a project. The ability to add small increments of capacity, often referred to as “modularity,” is another form of planning flexibility, as is the ability to construct a plant very rapidly to take advantage of current market conditions. Demand response, such as the response of aluminum smelters to wholesale price excursions, where it may take several months to efficiently shutdown or restart an industrial facility, is an example of long-term flexibility.

The value of planning flexibility was demonstrated during the energy crisis of 2000-2001. Few of the conventional power plants that entered construction during the crisis contributed to moderating the elevated prices that ensued, because the episode was over before construction was completed. What contributed most to re-establishing supply-demand equilibrium, instead, was reduced irrigation, reduced industrial and aluminum smelter load, and other demand response programs where consumers reduced demand in response to financial inducements. Other examples of the value of planning flexibility included reductions in spill that made additional hydro generation available and diesel generators that came on-line very quickly.

The treatment of flexibility, and in particular long-term planning flexibility, distinguishes the Council’s study and analytical technique from many of the techniques currently used to evaluate resource plans. This distinguishing feature is critical to the Council’s evaluation of risk. The following section describes how the portfolio model captures the value of modularity, short-lead time, and cost-effective deferral of construction.

Resource Additions and Decision Criteria

Planning flexibility allows a plan to accommodate changes from one future to another. By automating this process and applying probabilities to the various futures, the portfolio model can estimate the expected cost of accommodating the full range of futures. Plans containing resources that have more planning flexibility can manage this accommodation at a lower expected cost, other things being equal.

The value of flexibility stems from the ability to change plans when unforeseen events occur. This implies that a risk model must incorporate at least two special features.

First, a risk model must have the ability to add resource capacity without the benefit of perfect foresight. Most production cost or system simulation models capable of capacity expansion use techniques that assume perfect foresight. For example, these models may remove resources that do have sufficient value in the market to cover forward going fixed costs or add resources that would make a risk-adjusted profit in the market. An iterative process removes or adds resources until all new resources would just cover their risk-adjusted costs. Alternatively, a capacity expansion model may choose a capacity expansion schedule that minimizes cost. Both of these approaches must determine future hourly costs and prices to feed back to the capacity expansion algorithm. This feedback determines whether some adjustment to the construction schedule is necessary. If the model modifies the schedule, of course, the model must re-estimate future costs and prices change. The process repeats until the model finds a solution. These estimates of future costs and prices represent perfect foresight regarding how resources, costs, and prices affect one another. Perfect foresight, however, is contrary to the principles of risk analysis.[11]

Second, a risk model that incorporates capacity expansion must have a decision rule that determines whether to build or continue building. Because a risk model cannot use perfect foresight, the value of this criterion must use information about the current situation or about the past. Of course, different resources may use different criteria. A good test of a decision criterion, as it turns out, is whether it reduces cost and risk.

A decision criterion need not be perfect. The assessment of the value of planning flexibility relies on how well a resource plan performs when circumstances *do not* materialize as planned. As long as the decision criterion add resources and makes wrong forecasts (from the standpoint of perfect foresight) in a realistic manner, it could be deemed adequate.

The Council evaluated several approaches to decision criteria. For conventional thermal resources and wind generation, the approach that performed best incorporates information about resource-load balance and forward prices for fuel and electricity prices. Specifically, the model uses a three-year average of load growth and any change in resource capability to determine when in the future resource-load balance would cross below a given threshold. The selection of the threshold is itself part of the choice the model makes to minimize cost or risk. In each simulation period and for each resource candidate, the model determines whether the crossover point is less than the construction time required for that resource.

If the model needs a resource to meet anticipated future load, the criterion consults pertinent forward prices for each resource. For example, for a gas-fired power plant, the model would estimate the plant's value from forward prices for electricity and natural gas and compare those

[11] A peculiar side effect of perfect foresight models is they often lead decision makers to rely on the market. Capacity expansion models with perfect foresight add power plants precisely when they have greatest value. Following this approach, however, leads to market prices that match the fully allocated cost of the capacity expansion alternative or to long-term marginal expansion costs that match market prices. Given that the decision maker is no better building a plant than she would be if she purchased firm power in the market, there is little incentive to incur the considerable risks and challenges of building.

to capital and other fixed costs to determine whether the plant would pay for itself. If the plant would pay for itself, construction proceeds; if not, the model compares the value of the plant to that of alternatives. If the plant cannot pay for itself but is still the least expensive alternative, construction continues.

The model uses forward prices for electricity, natural gas, and other commodities, but it cannot use perfect foresight. Consequently, the model estimates forward prices using the assumption that futures and forward prices closely track current prices. This relationship is apparent in data for many commodities, including natural gas and electricity, where storage of the commodity is limited. The average commodity price over the last 18 months is the forecast of forward prices, reflecting the fact that it often takes a while for perceptions about long-term price to change.

Each resource that is a candidate for capacity expansion uses its decision criterion to control progress on construction, depending on where the resource is in its construction cycle. The decision criterion typically assumes one of two values, corresponding to either "Go" or "No Go" instruction, as illustrated in Figure 6-15.

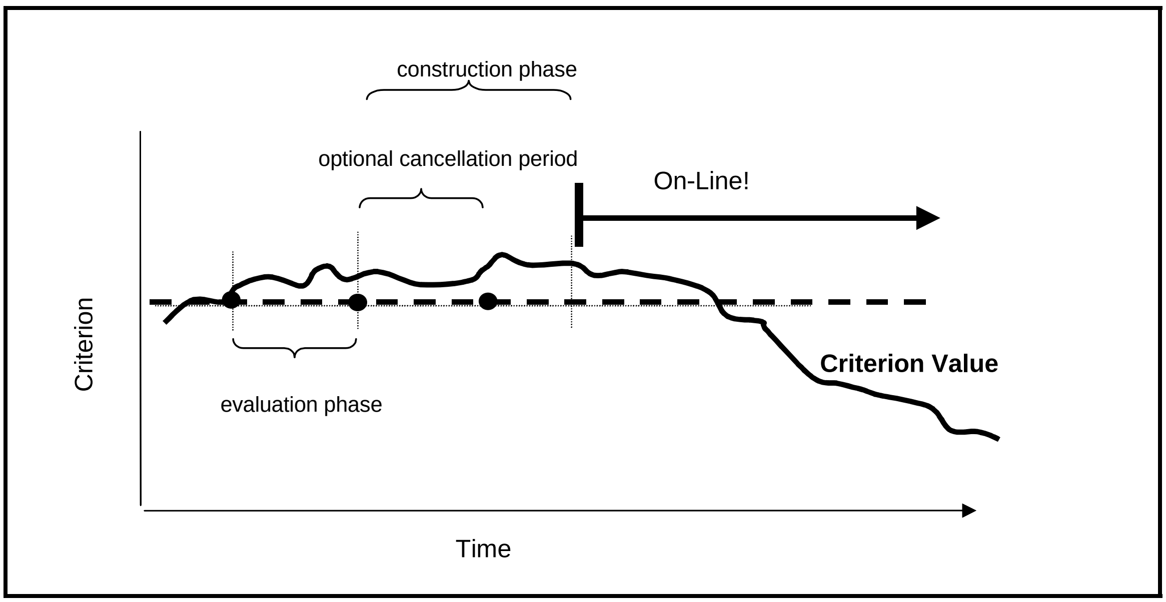

Figure 6-15: The Decision Criterion Value Over Time

The construction cycle for power plants typically consists of three distinct periods. (See Figure 6-16). During the first period, planning, siting, and permitting takes place. The regional portfolio model assumes planning costs are sunk. The purpose of the plan, in fact, is to determine for which, and for how much of each resources the region should complete such preliminaries. The second period commences with the first substantial, financial commitment. This might include activity such as substation and building construction, or an initial order for boilers or turbines. During this critical second period, the plant owner may delay or cancel construction if circumstances dictate. This period may last from several months to several years, depending on the resource, or it may not exist at all. The regional model captures this flexibility by delaying or canceling construction when the decision criterion indicates progress would not be advantageous. The model then incurs mothballing and cancellation costs for the plant. After this second period, however, a final commitment typically is required which compels the plant owner to finish construction. An example of an event that would trigger the third and final period is the receipt of and final payment for the turbine or boiler. These items are often the largest, single expense during construction. During this third period, construction activity in the model ignores the decision criterion.

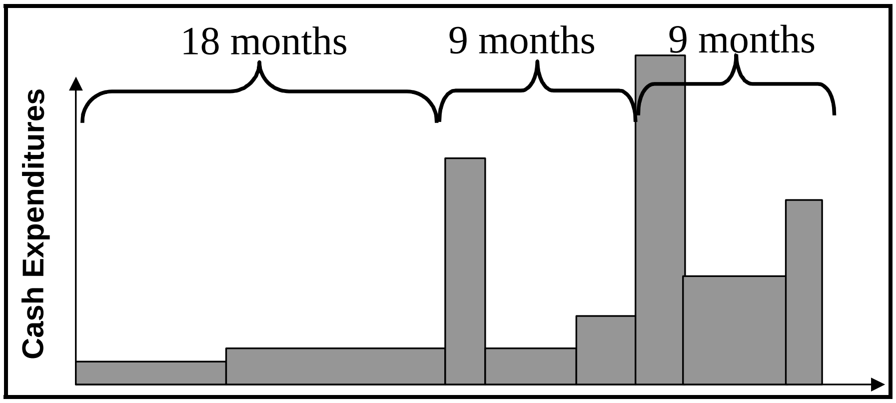

Figure 6-16: Stages of Cash Flow

Given how important the decision criterion is to assessing planning flexibility, it is natural to ask what alternatives exist and why the Council chose this particular decision rule. The first rule implemented in early versions of the portfolio model was valuation using forward prices, much as described above. One concern that arose when consideration turned to valuing conservation is that conservation often received value by virtue of “being there” when high market price excursions occurred. Resources that used only valuation in the market could only react to these excursions; often completing construction after the excursion subsided. Although this may help describe behavior during the 2000-2001 energy crisis, a more experienced market will probably pay careful attention to physical resource requirements in the future. Moreover, when a resource-load balance criterion replaced the market valuation criterion in the portfolio model, the feasibility space and its efficient frontier displayed reduced risk at no increase in cost. Resource-load balance does a better job of predicting the need for resources.

Resource-load balance alone, however, presents some problems as a decision criterion. An examination of particular futures revealed unrealistic behavior. Resource-load balance ignores economics completely. Given a future with high gas prices, for example, the portfolio model would be as likely to develop a gas-fired turbine as a coal plant if it has a choice between the two. Consequently, the criterion in the final version of the portfolio model gives consideration first to resource-load balance and then uses plant valuation to make the resource choice.

Conservation uses a slightly different decision criterion. Conservation can introduce thorny problems, like cost shifting for ratepayers and revenue recovery for load-serving entities. Consequently, special regulatory or administrative intervention is typically necessary. Cost effectiveness has been the standard that administrators use to deem the type and amount of conservation to pursue.

Because conservation uses a cost effectiveness standard, a criterion that resembles such a standard seems appropriate. However, the challenges in constructing a cost effectiveness criterion are several.

- Cost effectiveness levels change over time as market prices for electricity change, although administrators tend to base them on long-term equilibrium prices for electricity. Models that estimate equilibrium prices for electricity are sensitive to commodities that have been less volatile than electricity prices, such as natural gas price. Regardless, cost effectiveness standards are subject to uncertainty and change depending on the particular future.

- Because they are often determined administratively, they change more slowly than commodity prices. Moreover, the time between changes in efficiency standards and when the conservation measure starts to contribute can be a year or more, while load-serving entities develop their budgets and ramp up programs. Thus, there is considerable lag time between changes in commodity prices and changes in conservation energy rate of addition.

- Some types of conservation become institutionalized, such as that associated with new codes and standards for building construction. Once the codes pass into law, the corresponding measures are no longer directly subject to the cost effectiveness standard. Thus, the decision criterion for this kind of conservation is "sticky downward." It does not decrease, and it increases only when the cost effectiveness standard passes the previous "high-water mark."

- The NW Power Act requires that the power plan assign a ten percent cost advantage to the acquisition of conservation. By using a criterion that accessed the supply curve as a level at least 10 percent higher than a market-based cost effectiveness standard, the portfolio would accommodate this requirement.

- A long-standing Council objective has been to understand what value there may be in sustained, orderly development of conservation. Is there any advantage to this policy over the sustained, orderly development of any other resource? Is there any cost or risk advantage to developing more conservation than a conventional cost effectiveness standard would suggest?

These considerations drove the design of the decision criteria for conservation. In the case of conservation, the decision criterion takes the form of a price. This price and a supply curve determine how much conservation to develop in a given period. Both lost-opportunity and discretionary conservation[12] criteria are the sum of two terms. The first term approximates the cost-effectiveness standard. This is a "myopic" estimate of cost effectiveness, which depends on the specific future and changes over time in that future. The second term determines how much additional conservation to deploy compared to the cost effectiveness level. This second term, a price adjustment, is under the control of the logic that helps the portfolio model find the least-cost plan, given a fixed level of risk. (See Figure 6-7.)

The specific rules for estimating the going-forward cost effectiveness standards appear in Appendix L. The discussion of conservation in Chapter 3 shows the effect of alternative rules on

[12] The description of these classes of conservation appears in Chapter 3

cost, risk, and acquisition levels. The reader will also find in Appendix L a discussion of the effect that the shape of the supply curve has on the value of conservation under uncertainty. This discussion explains the relation between the price adder and the reduction of cost.

Finally, this section has emphasized the role of planning flexibility and the decision criteria, but the reader should remember there is another important element that determines construction. The logic that helps the portfolio model find the least-cost plan plays an equal, if not larger role in which resources can show up in a given future. This logic determines which resources have completed the pre-construction stage and, therefore, which resources are *available* for construction.

This is the purpose of the resource plan produced by the portfolio model: to determine which resources to prepare and when to commit to their deployment. The next chapter describes the resource plan that the Council selected and provides additional interpretation of the plan's schedule for construction and action.